arXiv:1406.6095v1 [astro-ph.CO] 23 Jun 2014

DISS. ETH NO. 20642

The information content of galaxy surveys

A dissertation submitted to

ETH ZURICH

for the degree of

Doctor of Sciences

presented by

Julien Carron

MSc ETH in Physics

born on Dec 26th, 1985

citizen of Sion

accepted on the recommendation of

Prof. Dr. Simon Lilly, examiner
Dr. Adam Amara, co-examiner
Prof. Dr. Luca Amendola, co-examiner

2012

Contents

Contents vi

List of Figures viii

List of Tables ix

Abstract xi

Résumé xiii

1 Introduction 1

1.1 Stochasticity in cosmological observables 2

1.2 Fisher information for cosmology : a first look 3

1.3 N-point functions 6

1.3.1 N-point functions from the density : characteristic functional 6

1.3.2 The density from N-point functions : determinacy of the moment problem 9

1.3.3 N-point functions from discrete populations, poisson samples 9

1.4 Gaussian and non-Gaussian matter density fields 11

1.5 Structure of the thesis 15

I Quantifying information **17**

2 Shannon entropy and Fisher information **19**

2.1 Fisher information and the information inequality 20

2.1.1 The case of a single observable 21

2.1.2 The general case 23

2.1.3 Resolving the density function: Fisher information density 24

2.2 Jaynes Maximal Entropy Principle 26

2.2.1 Information in maximal entropy distributions 29

2.2.2 Redundant observables 31

2.2.3 The entropy and Fisher information content of Gaussian homogeneous fields 32

2.3 Appendix 36

2.3.1 Measures of information on a parameter 36

2.3.2 Cramér-Rao bound and maximal entropy distributions 39

3 Information within N-point moments **43**

3.1 One variable 44

3.1.1 The Christoffel-Darboux kernel as information on the density itself . 47

3.2 Several variables 48

3.2.1 Independent variables and uncorrelated fiducial 50

3.2.2 On the impact of noise 51

3.2.3 Poissonian discreteness effects 53

3.3 Some exactly solvable models for moment determinate densities 53

3.3.1 Moment indeterminate densities 60

3.4 Appendix 66
3.4.1 Derivation for the beta family 66
3.4.2 General hierarchical systems, recursion relations. 70

II Applications for cosmology **73**

4 On the combination of shears, magnification and flexion fields **75**
4.1 Introduction 75
4.2 Linear tracers 76
4.2.1 Linear tracers at the two-point level 79
4.2.2 Joint entropy and information content 79
4.2.3 Weak lensing probes 81
4.3 Results 83
4.4 Appendix 86

5 On the use of Gaussian likelihoods for power spectra estimators **91**
5.1 Introduction 91
5.2 One field, gamma distribution 92
5.3 Several fields, Wishart distribution 95
5.4 Summary and conclusions 98

6 N-point functions in lognormal density fields **101**
6.1 Introduction 101
6.1.1 Notation and conventions 103
6.1.2 Definition and basic properties of correlated lognormal variables . . 103
6.2 Fundamental limitations of the lognormal assumption 104
6.2.1 Largest scales 105

6.2.2 Smallest scales . . . 106
6.3 Explicit fields with the same hierarchy of N-point functions . . . 108
6.3.1 The problem with tailed fields . . . 108
6.3.2 Continuous family . . . 111
6.3.3 Discrete family . . . 113
6.3.4 Appendix . . . 114
6.4 Information at all orders for the one dimensional distribution . . . 115
6.4.1 Lack of information in the moments . . . 118
6.4.2 A q-analog of the logarithm . . . 119
6.4.3 Comparison to standard perturbation theory . . . 122
6.4.4 Appendix . . . 123
6.5 Connection to N-body simulations . . . 126
6.6 Information at all orders for uncorrelated fiducial . . . 130
6.6.1 Cumulative efficiencies . . . 132
6.6.2 Impact of discreteness effects . . . 133
6.6.3 The reason for a roughly parameter independent factor of improvement . . . 137
6.7 The general case : a difficult Lagrange interpolation problem . . . 139

7 Information escaping the correlation hierarchy of the convergence field 141

7.1 Introduction . . . 141
7.2 Fisher information coefficients . . . 143
7.3 Restoration of the information . . . 146
7.4 Conclusions . . . 147

Summary of main results and outlook 149

Bibliography 162

List of Figures

1.1 The shape of the one-point distribution function of the matter density fluctuation field according to second order perturbation theory. 12

1.2 Evolution of the matter density field distribution with cosmic time in a N-body simulation. Image courtesy of A. Pillepich 14

3.1 The cumulative efficiencies of the moments of the lognormal distribution to capture its information content 62

3.2 The cumulative efficiencies of the moments of the Weibull distribution to capture its information content 64

3.3 The cumulative efficiencies of the moments of the stretched exponential function to capture its information content 65

4.1 The ratio of the effective noise when including flexion and magnification fields to the level of noise considering the shears only, as a function of the angular scale. 84

4.2 The increase of the dark energy FOM including size information as well as flexion information, over the shear only analysis, as a function of the maximal angular multipole. 87

6.1 The one dimensional chain of points used to test the positivity of the spectrum of the A field, and the apparition modes with negative spectrum on small scales, indicating that the lognormal field is no longer a well defined assumption for a ΛCDM cosmology. 106

6.2 Three density functions with the same moments as those of the lognormal distribution at all orders. 112

6.3 The fraction of the total information content that is accessible through the extraction of the entire series of moments of the lognormal field, as function of the square root of the variance of the fluctuations. 120

6.4 Distribution of the information within the moments of the lognormal distribution 120

6.5 The information density of the lognormal distribution and its approximation through the associated orthonormal polynomials, for fluctuations of unit variance 121

6.6 Comparison of the information content of the first moments of the fluctuation field between standard perturbation theory and the lognormal model . 124

6.7 Comparison of various estimates of the error bar on the linear power spectrum amplitude, constrained using the power spectra of the overdensity field or of the log-density field in an N-body simulation. 127

6.8 The fraction of the total information content that is accessible through extraction of the full series of N-point moments of the lognormal field for an uncorrelated fiducial model, for a parameter entering the two-point function at non zero lag, but not the variance. 133

6.9 Discreteness effects destroying the information contained in the underdense regions of the lognormal distribution 135

6.10 The information content of the first N-point moments in a Poisson sample of a lognormal field, as function of the mean number of objects per cell, normalised to the information content of the lognormal field. 136

6.11 Ratios measuring the sensibility of the spectrum to that of the variance for different cosmological parameters 138

7.1 The three parameters entering the generalized lognormal model, as function of the variance of the effective matter density fluctuations 144

7.2 The cumulative efficiencies of the moments of the convergence in capturing Fisher information as function of the variance of the effective matter density fluctuations. 146

7.3 The efficiency $\epsilon_{N=10}$ of the first 10 moments to capture Fisher information, for the convergence field and lognormal field, as well as for the logarithmic transform of the convergence field. 147

7.4 The Fisher information density of the lognormal and the convergence field. 148

7.5 The cumulative information on the amplitude of the spectrum within either the spectrum or two-point function of the one dimensional lognormal field. . 151

List of Tables

4.1 Dispersion parameters for the shear, flexion and magnification fields used in our numerical analysis. 83

4.2 Ratio of the marginalised constraints considering size information as well as flexion to that of shear only analysis. 86

6.1 The factors of improvements in constraining power of the spectrum of the log-density field in the lognormal model and as seen in simulations 130

Abstract

This research is a contribution to our understanding of the information content of the cosmological dark matter density field, and of the means to extract this information. These questions are of prime importance in order to reach closer for solutions to current fundamental issues in cosmology, such as the nature of dark matter and dark energy, that future large galaxy surveys are aiming at. The focus is on a traditional class of observables, the N-point functions, that we approach with known information theoretic tools, Fisher information and Shannon information entropy. It is built out of two main parts, the first presenting in details the mathematical methods we used and introduced, and the second the cosmological research that was performed with these tools.

A large fraction of this thesis is dedicated to the study of the information content of random fields with heavy tails, in particular the lognormal field, a model for the matter density fluctuation field. It is well known that in the nonlinear regime of structure formation, the matter fluctuation field develops such large tails. It has also been suggested that fields with large tails are not necessarily well described by the hierarchy of N-point functions. In this thesis, we are able to make this last statement precise and with the help of the lognormal model to quantify precisely its implications for inference on cosmological parameters : we find as our main result that only a tiny fraction of the total Fisher information of the field is still contained in the hierarchy of N-point moments in the nonlinear regime, rendering parameter inference from such moments very inefficient. We show that the hierarchy fails to capture the information that is contained in the underdense regions, which at the same time are found to be the most rich in information. We find further our results to be very consistent with numerical analysis using N-body simulations. We also discuss these issues with the help of explicit families of fields with the same hierarchy of N-point moments defined in this work. A similar analysis is then applied to the convergence field, the weighted projection of the matter density fluctuation field along the line of sight, with similar conclusions. We also show how simple mappings can correct for this inadequacy, consistently with previous findings in the literature.

These results were made possible using an expansion of the Fisher information matrix in uncorrelated components associated to N-points moments of successive orders. An entire chapter is dedicated to this expansion, investigating its properties and making a connection to the moment problem in the field of mathematics. Some simple models exactly solvable at all orders are also presented.

Beside these investigations of the statistical power of the hierarchy of N-point moments, we also study the combination of various probes of the convergence field, including the magnification, shear and flexion fields, in particular at the two-point level. We use Shannon information entropy to discuss the simple structure of the information within these tracers of the lensing potential field. We then evaluate the prospects for such a combination according to current understanding of the relevant dispersion parameters. Finally, we revisit known derivations of the Fisher information matrix for Gaussian variables, commenting in this light on the use of Gaussian likelihoods for power spectra or two-point correlation function estimators in cosmology. We point towards the fact that despite their motivation from the central limit theorem, care must be taken in the case of a large number of fields, as this assumption assigns too much information to the observables.

Résumé

Cette thèse est une contribution à notre compréhension de l'information utile à des fins cosmologiques contenue dans le champ de matière noire de notre Univers, ainsi que des procédés pour extraire cette information. Ces questions sont essentielles dans l'optique de s'approcher d'éléments de réponses à des questions fondamentales de la cosmologie moderne, telles que la nature de la matière noire et de l'énergie sombre. Dans ce travail, nous concentrons principalement nos efforts sur une classe traditionnelle d'observables, c'est à dire la hiérarchie des fonctions à N points, que nous approchons en utilisant deux outils empruntés à la théorie de l'information : l'information de Fisher et l'entropie de Shannon. Cette thèse est constituée de deux parties principales, la première présentant et développant ces méthodes mathématiques pour notre fin, la deuxième la recherche en cosmologie proprement dite qui a été effectuée avec ces outils.

Une portion importante de ce travail est dédiée à l'étude de l'information de Fisher contenue dans les champs aléatoires avec forte asymétrie droite, et en particulier dans les champs log-normaux. Il est bien connu que dans le régime non linéaire de la formation des structures, le champ de densité de matière noire développe une telle asymétrie, et il a été également suggéré que cette asymétrie pénalise les fonctions à N points dans leur capacité à décrire ces champs. Dans cette thèse, nous formulons ce dernier aspect plus précisément, et calculons dans plusieurs modèles son impact sur notre capacité à extraire les paramètres cosmologiques de ces observables. Nous trouvons qu'une petite fraction seulement d'information reste accessible à la hiérarchie des fonctions à N points dans le régime non linéaire. Nous montrons que ces observables sont inadéquates à capturer l'information du champ dans les régions sous-denses, qui elles-mêmes contiennent la plus grande part de l'information dans ce régime, et confrontons avec succès ces résultats à des simulations numériques à N corps. Nous discutons également plusieurs de ces aspects avec l'aide de quelques exemples explicites de champs aléatoires qui possèdent les mêmes fonctions à N points à tous les ordres. Nous effectuons une analyse similaire pour le champ de convergence avec des résultats inchangés. Nous montrons comment de simples

transformations non linéaire permettent de corriger ces problèmes, de manière comparable à d'autres résultats déjà présents dans la littérature.

Ces résultats sont rendus possibles par une expansion de la matrice d'information de Fisher en composantes associées de manière univoque aux membres successifs de la hiérarchie des moments à N points que nous introduisons dans ce travail. Un chapitre entier est consacré à cette expansion, où nous discutons ses propriétés principales, ainsi que les liens avec le problème des moments en mathématique. Nous résolvons aussi quelques modèles suffisamment simples pour permettre la dérivation analytique d'une solution exacte à tous les ordres.

Nous présentons également une étude sur la combinaison de différents traceurs du champ de convergence. Nous considérons les champs de grandissement, de cisaillement faible et de flexion, tous directement reliés à la convergence, notamment leurs fonctions à deux points, et utilisons l'entropie de Shannon pour discuter leur information jointe. Nous évaluons les avantages d'une analyse jointe de ces traceurs selon notre compréhension actuelle des paramètres de dispersion. Finalement, nous revoyons des dérivations connues de la matrices de Fisher pour des champs Gaussiens, ce qui nous permet de commenter l'usage fréquent de statistiques Gaussiennes pour des estimateurs de spectres de puissance ou de fonctions à deux points en cosmologie. Bien que la forme Gaussienne soit motivée par le théorème de la limite centrale, nous notons que cette hypothèse assigne trop d'information aux spectres dans le cas d'une analyse jointe de plusieurs champs aléatoires.

1 Introduction

After the cosmic microwave background radiation, the study of the formation and evolution of the structures on the large scales of our Universe forms one of the pillars of modern cosmology. These structures can be mapped by galaxy surveys, and cosmological observables derived from these surveys such as the galaxy two-point correlation function or its Fourier transform the power spectrum are central to the field, and are used to contrast predictions of cosmological models to observations (Percival et al., 2010; Pope et al., 2004; Tegmark et al., 2006). Despite essential successes in the last two decades with the emergence an observationally very successful concordance cosmology, this description of our Universe, the ΛCDM model, is still very mysterious, with only about 5% of the energy budget of the Universe being made of the matter which we have daily experience of (Komatsu et al., 2011). The real nature of the two dominating components, the dark energy and the dark matter, remains unclear to this day and is the heart of a large scientific effort. Both the dark matter density field as well as the impact of dark energy on the geometry of the Universe can now in principle both be observed with the help of weak lensing (Bartelmann and Schneider, 2001; Schneider et al., 2006). It is thus believed that large galaxy surveys able to reach for the lensing signal are going to play an increasingly important role towards these fundamental issues in cosmology (Albrecht et al., 2006).

Of course, in order to assess some set of observables as valuable for cosmology, and to design an experiment towards its extraction, it is essential to understand both our capabilities to extract it, as well as the robustness and pertinence of the predictions of our model. These aspects, often of statistical nature, are the very backbones of this thesis. Our main aim in the present research was to try and quantify these aspects in several situations relevant to cosmology, focussing on the dark matter density field, or its weighted projection along the line of sight, the weak lensing convergence field, contributing in this way to our understanding of the information content of galaxy surveys. We review in this introductory chapter the known tools that we have built upon as well as the class of observables we have focused on, putting thus our work in context.

1.1 Stochasticity in cosmological observables

All major predictions and measurements that are used to test our understanding of our cosmological model are meaningful only in a statistical sense. Indeed, our inability to observe initial conditions, which we may tentatively evolve, as well as the complexity of some of the physical processes involved render in general a statistical description unavoidable. For this reason, a key element that determine to an often decisive extent what observable will be of interest for the purpose of the analyst is the probability density for the realisation of the fields from which the observables are derived. Typically a CMB temperature map, or a galaxy density field, from which one measures for instance the two-point correlation function. This element of stochasticity is sometimes referred to as *cosmic variance*, a denomination that we adopt in the following. One must generically include other sort of stochasticity on top of the cosmic variance, that we refer to as *noise*, for instance due to the specificities of the instrumentation, filling another gap between model predictions and actual data outputs.

We need to introduce some notation :

We always write a probability density with p, at times adding a subscript indicating to which random variable it refers to for clarity. In the case of cosmological fields, these probability densities are generically high dimensional, describing the joint occurrence of fields values at different points. Typically, when the random variable are the values $(\phi_1, \cdots, \phi_d)$ of a a field ϕ at points $(x_1, \cdots, x_d)$, then p is a function of d variables, a d-point probability density function. The position label x itself can have various meanings in diverse cosmologically relevant situation. It can have for instance dimension $n = 1$ (Lyman-α forest), $n = 2$ (weak lensing tomography, projected density fields, CMB), or $n = 3$ (redshift surveys).

The joint density for the realisation of the field ϕ at all points can be written conveniently as the functional $p[\phi]$. Expectation values of observables f are given formally as

$$\langle f \rangle = \int \mathcal{D}\phi \, p[\phi] f[\phi], \tag{1.1}$$

an infinite dimensional integral. It should be kept in mind that such probability densities $p[\phi]$ are however not always very well defined and intrinsically difficult to handle, except in some cases. Expectation values 1.1 can nevertheless be understood as the limit of a finite dimensional, well defined average

$$\langle f \rangle = \int d^d\phi \, p\left(\phi_1, \cdots, \phi_d\right) f\left(\phi_1, \cdots, \phi_d\right) \tag{1.2}$$

over a finite sample of the field, with large d. In a harmless abuse of terminology we may identify at times in this work such finite samples of the field with the field itself, especially when dealing with N-body simulations, that have of course only a finite number of spatial resolution elements.

Homogeneity, isotropy, ergodicity

Cosmic variance in the sense defined above is the stochasticity of the data due to the fact that we observe one particular realisation of a random field, namely that of our own Universe (or of the observed part of the Universe, in which case one can also refer to a component of sample variance). It is a fundamental limitation in the sense that this variability can never be beaten down, as this would ultimately require the observation of several universes governed by the same density functions, which is a mathematical construct useless to our purposes.

Within this framework, one relies on several assumptions, namely that of statistical homogeneity, isotropy and ergodicity. The first two express the absence of preferred locations and directions in the Universe. Mathematically speaking, all density functions are required to be invariant under spatial translations and rotations,

$$p\left(\phi(x_1), \cdots, \phi(x_d)\right) = p\left(\phi(x_1 + r), \cdots, \phi(x_d + r)\right) = p\left(\phi(R \cdot x_1), \cdots, \phi(R \cdot x_d)\right), \quad (1.3)$$

for any translation vector r and rotation matrix R. These two important assumptions can be tested and are confronted to observations. Of course, homogeneity and isotropy do not apply to fields in redshift space coordinates. The third, ergodicity, states that we can reinterpret the ensemble averages in equation 1.1 to be spatial averages. We expect this assumption to be correct as long as the spatial averages can be made over sufficiently large volumes, or using widely separated samples, assuming that correlations at large distances decays quickly enough to zero. Under these conditions, so called ergodic theorems can indeed be proven. However, this assumption cannot be fundamentally tested and we have no choice but to take it as granted in order to obtain useful results out of this mathematical approach.

Very often of primary interest are the zero mean, dimensionless fluctuations δ of ϕ, defined as

$$\delta(x) = \frac{\phi(x) - \bar{\phi}}{\bar{\phi}}, \quad (1.4)$$

where $\bar{\phi} = \langle \phi(x) \rangle$ is the mean of the field, independent of position x by homogeneity.

1.2 Fisher information for cosmology : a first look

Inference on model parameters might appear extremely simple in principle. For a set of model parameters $\boldsymbol{\theta} = (\alpha, \beta, \cdots)$ of interest, and the observed field ϕ, probability theory tells us that we must update our knowledge of $\boldsymbol{\theta}$ with the simple rule,

$$p(\boldsymbol{\theta}|\phi) = \frac{p(\phi|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int d\boldsymbol{\theta}\, p(\phi|\boldsymbol{\theta})p(\boldsymbol{\theta})}. \quad (1.5)$$

In this equation, the density $p(\boldsymbol{\theta})$ describes our prior state of knowledge on $\boldsymbol{\theta}$, and $p(\phi|\boldsymbol{\theta})$, viewed as function of the parameter is called the likelihood. On the lefthand side, $p(\boldsymbol{\theta}|\phi)$

is called the posterior. Of course, the simplicity of this formula should not hide the very high complexity of its implementation for typical cosmological instances. In particular, the likelihoods $p(\phi|\boldsymbol{\theta})$ are in general only poorly known, and the very high dimensionality of this object requires the compression of ϕ to some smaller subset of observables, whose statistics are set by the likelihood, all carrying some of the amount of the information that the likelihood carried originally. It is thus clearly of the uttermost interest to be able to quantify more precisely this information, both that of the original likelihood as well as that of the different observables. This is where Fisher information comes into play.

It seems fair to say that the use of Fisher information in cosmology begins, though indirectly, with Jungman et al. (1996a,b), two works in the context of CMB experiments aiming at measuring the temperature fluctuation spectrum C_l. In these works, it is argued that the posterior for the parameters will be approximately Gaussian. Let us consider for simplicity the case of a single parameter of interest, α, as the discussion or an arbitrary number of parameters holds essentially unchanged. With the true value, or best fit value of the parameter defined as α_0, the posterior is assumed to have the form

$$p(\alpha|\{C_l\}) \propto \exp\left(-\frac{1}{2}(\alpha-\alpha_0)^2 i_\alpha(\alpha_0)\right), \tag{1.6}$$

with the number i_α is defined as

$$i_\alpha(\alpha_0) = \sum_l \frac{1}{\sigma_l^2}\left(\frac{\partial C_l(\alpha_0)}{\partial\alpha}\right)^2. \tag{1.7}$$

In this equation, σ_l^2 is the variance of the estimates of C_l, including cosmic variance, incomplete sky coverage and detector noise.

Under the assumption (1.6), it is clear that $1/i_\alpha$ is the variance σ_α^2 of the parameter. Very interestingly, from its definition (1.7) we see that this variance can be evaluated prior obtaining data, if a reasonable fiducial point α_0 can be chosen and the model predictions of the spectrum are given. Provided the assumptions made there are correct, this is making the approach of Jungman et al. (1996a,b) quite powerful, providing us with a rather great understanding of the capabilities of the experiment.

It is worthwhile spending a bit more thoughts on i_α defined in (1.7). It is a special case of an expression that weights the derivatives of some set of observables O_i according to their covariance matrix Σ_{ij},

$$i_\alpha(\alpha_0) = \sum_{i,j} \frac{\partial O_i(\alpha_0)}{\partial\alpha}\left[\Sigma^{-1}\right]_{ij}\frac{\partial O_j(\alpha_0)}{\partial\alpha}. \tag{1.8}$$

We recover (1.7) by setting the observables to be the spectrum, and the covariance matrix to be diagonal, as required for a perfectly Gaussian CMB map. The number (1.8) has an array of fundamental properties, none of them being difficult to show :

- It is a non negative number, that becomes larger for a smaller covariance matrix or a larger impact of α on the observable, and vice versa.

- i_α corresponding to independent observables (i.e. with no covariance) is simply the sum of their respective i_α.
- Adding an observable O_{n+1} to a set $(O_1, \cdots O_n)$ can only increase i_α, and not decrease it.
- i_α is identical to the expected curvature of a least squares fit to the observables with the given covariance matrix[1], provided α_0 is the parameter value that gives the least squared residuals.

These properties are very consistent with what we would expect from a measure of information on the parameter α and are making the number (1.8) a promising candidate for such a measure. However, it is clearly not the end of the story, since it depends only on the chosen set of observables and their covariance matrix, but neglects all other aspects of the probability density $p(\phi|\alpha)$. The link with Fisher information, a well known tool in statistics, was then exposed and extended to other areas of cosmology in works such as Tegmark (1997); Tegmark et al. (1997). It was noted that the number (1.7) is identical to

$$I(\alpha) := -\left\langle \frac{\partial^2 \ln p}{\partial \alpha^2} \right\rangle, \tag{1.9}$$

where p is a Gaussian likelihood for the noisy CMB temperature fluctuation field. The connection between equation (1.9) and covariances on parameters was also used in an astrophysical context earlier in Amendola (1996). Equation 1.9 is the Fisher information in p on α, a most sensible measure of information on parameters, whose properties and link to (1.8) we will have the occasion to discuss extensively. For several parameters, it becomes the Fisher information matrix

$$F_{\alpha\beta} := -\left\langle \frac{\partial^2 \ln p}{\partial \alpha \partial \beta} \right\rangle, \tag{1.10}$$

Since then, such Fisher information matrices for Gaussian variables and the assumption (1.6) have been used routinely in cosmology in order to assess the capabilities of some future experiments.

Two comments are in order at this point :

First, the definition (1.10) of the Fisher information matrix is the most common in cosmology. However, in this thesis, we will rather use the alternative

$$F_{\alpha\beta} := \left\langle \frac{\partial \ln p}{\partial \alpha} \frac{\partial \ln p}{\partial \beta} \right\rangle \tag{1.11}$$

as the definition of the information matrix. These two definitions can be shown to be equivalent for any probability density function using the fact that probability densities are normalised to unity, $0 = \frac{\partial}{\partial \alpha} \langle 1 \rangle = \langle \partial_\alpha \ln p \rangle$. While (1.10) conveniently presents the

[1] It should be noted that this identification to a curvature in a least square fitting procedure holds only if the covariance matrix is treated as parameter independent.

information matrix as a curvature matrix, it will become clear in chapter 3 that (1.11), making a reference to the *score function* $\partial_\alpha \ln p$, is in fact much more fundamental for our purposes. This form generalises more easily to non normalised density functions as well.

Second, as discussed above, Fisher information is often interpreted in cosmology as an approximation to the parameter posterior, approximated as a Gaussian with covariance matrix F^{-1}. The Gausisan approximation, as well as the identification of the Fisher information matrix with the inverse covariance matrix are of course only assumptions that can fail, at times severely. This is especially true when marginalising within this approach over poorly constrained parameters, whose distribution often cannot be approximated by a Gaussian shape (see for example Wolz et al. (2012)), giving rise to results that are difficult to interpret. In this thesis the focus is on the more orthodox interpretation of the Fisher information matrix as a very meaningful and well defined measure of information, and not as an approximation to a posterior. In particular, we are not going to inverse the Fisher matrix or marginalise over a set of parameters, except in some instances making connections to results in the literature.

1.3 N-point functions

A very common class of observables, at the heart of this thesis, are the N-point functions, that we review briefly in this section. They are very convenient at least for two reasons. First, for Gaussian fields the mean and two-point function do contain the entire information in the field : in the language of orthodox statistics, they form a set of sufficient statistics. Second, the measurement of a (connected) three-point or higher order point function directly tests for non Gaussianity of the field.

1.3.1 N-point functions from the density : characteristic functional

In cosmology, the N-point function ξ_N is defined as the connected part of the N-point moment of the fluctuation field. These are most easily defined using the generating function technology, ubiquitous in any field theory. Consider first an arbitrary N-point moment of the field,

$$\langle \delta(x_1) \cdots \delta(x_N) \rangle \,. \tag{1.12}$$

We can write it, at least formally, as a derivative of the generating functional Z, or *characteristic functional*, essentially the Fourier transform of the density :

$$\langle \delta(x_1) \cdots \delta(x_N) \rangle = \frac{1}{i^N} \frac{\partial^N}{\partial J(x_1) \cdots \partial J(x_N)} Z[J] \bigg|_{J=0} , \tag{1.13}$$

with

$$Z[J] := \left\langle \exp\left(i \int d^n x \, J(x) \delta(x) \right) \right\rangle . \tag{1.14}$$

In other words, the N-point moments can be considered as the successive terms in an expansion of the generating functional in a power series in J. Note that there are cases, such as for instance the lognormal field, where the generating functional cannot be written as a power series, even for J very close to zero. In this case, the series should be considered as a formal power series regardless of convergence. The connected N-point correlation functions are then defined as the successive terms in the formal expansion of $\ln Z$:

$$\xi_N(x_1, \cdots, x_N) = \frac{1}{i^N} \frac{\partial^N}{\partial J(x_1) \cdots \partial J(x_N)} \ln Z \bigg|_{J=0} . \tag{1.15}$$

If all arguments x_1 to x_N are identical, these connected point functions become the familiar cumulants of the one dimensional density function $p(\phi(x))$.

The connected point functions are convenient since they are additive for uncorrelated fields. Indeed, if two fields are uncorrelated, then one finds directly from its very definition (1.14) that the characteristic function Z of the joint density is the product of the characteristic functions of each of the densities. Taking the logarithm and using the definition (1.15) shows that the connected functions just add up. From these relations (1.13) and (1.15) one can infer recursion relations for the connected point functions, as well as convenient diagrammatic representations, Feynman diagrams alike, where connected point functions are represented by connected graphs (Bernardeau et al., 2002; Szapudi, 2005, e.g.).

It holds that the very first connected point functions are identical to the first moments of the delta field,

$$\begin{aligned} \xi_2(x_1, x_2) \equiv \xi(x_1, x_2) &= \langle \delta(x_1)\delta(x_2) \rangle \\ \xi_3(x_1, x_2, x_3) &= \langle \delta(x_1)\delta(x_2)\delta(x_3) \rangle , \end{aligned} \tag{1.16}$$

but this is not the case anymore for higher N.

From homogeneity and isotropy, these functions are invariant under translations and rotations. In particular the two-point function is a function of a single argument,

$$\xi(x_1, x_2) = \xi \left(|x_1 - x_2| \right) . \tag{1.17}$$

Translation invariance allows conveniently the use of a description in terms of harmonics. In Cartesian space, with Fourier transform

$$\tilde{\delta}(\mathbf{k}) = \int d^n x \; \delta(x) e^{-i\mathbf{k}\cdot x}, \tag{1.18}$$

we have for any statistically homogeneous field the simple relation

$$\left\langle \delta(\mathbf{k}) \delta^*(\mathbf{k}') \right\rangle = (2\pi)^n \, \delta^D(\mathbf{k} - \mathbf{k}') P(\mathbf{k}), \tag{1.19}$$

where δ^D is the Dirac δ function and $P(\mathbf{k})$, the power spectrum, is the Fourier transform of the two-point function

$$P(\mathbf{k}) = \int d^n x \; \xi(x) e^{-i\mathbf{k}\cdot x}. \tag{1.20}$$

Is the field further statistically isotropic, the spectrum is only a function of the modulus k of the wavenumber. Similarly, one can define higher order spectra, the polyspectra, prominently the bispectrum for $N = 3$ and trispectrum for $N = 4$ through the Fourier transforms of the connected N-point functions, or equivalently the expectation of products of the Fourier modes of the field.

In this thesis, the distinction between connected and disconnected point functions, or the use of polyspectra rather than the N-point functions, are of no fundamental relevance, as they provide equivalent descriptions of the same source of information. We will not make a difference between a connected or disconnected point function. We regard a generic N-point moment

$$\langle \phi(x_1) \cdots \phi(x_N) \rangle \tag{1.21}$$

as a N-point function.

The prime example of a homogeneous isotropic random fluctuation field is of course the Gaussian field. Gaussian fields are very convenient for many reasons. They are stable under any linear transformations, such as smoothing, and also under convolutions. The celebrated central limit theorem states that sums of a large number of independent variables tend to have a Gaussian distribution under fairly generic conditions. Besides, they also arise as fields of maximum information entropy for a given two-point function. The Gaussian field is defined through

$$\ln p[\delta] = -\frac{1}{2} \int d^n x \int d^n y \, \delta(x) \xi^{-1}(x-y) \delta(y) + \text{cst}, \tag{1.22}$$

or, in Fourier space,

$$\ln p[\delta] = -\frac{1}{2} \int \frac{d^n k}{(2\pi)^n} \frac{|\tilde{\delta}(\mathbf{k})|^2}{P(\mathbf{k})} + \text{cst}. \tag{1.23}$$

The second representation shows that the Fourier modes of such a field are independent complex Gaussian variables with the correlations as given in (1.19). All finite d-dimensional joint densities are d-dimensional multivariate Gaussian distributions.

The characteristic functional can be evaluated in closed form. It is a standard result called the Gaussian integral.

$$Z[J] = \exp\left(-\frac{1}{2} \int d^n x \int d^n y J(x) \xi(x-y) \, J(y) \right) = \exp\left(-\frac{1}{2} \int \frac{d^n k}{(2\pi)^n} P(\mathbf{k}) |\tilde{J}(\mathbf{k})|^2 \right). \tag{1.24}$$

It follows that $\ln Z$ is a polynomial second order in J. It is then immediate that the connected point functions of the Gaussian field vanish for $N > 2$, since the derivatives of that order do vanish.

1.3.2 The density from N-point functions : determinacy of the moment problem

A key to several results of this thesis is the so-called *moment problem* and its *determinacy*. These are respectively the problem of finding a density given the hierarchy of N-point moments, and the question of whether a solution is unique or not. While not part of the usual cosmological literature, this topic is a well known area of research of mathematics, in particular for one dimensional densities (Akhiezer, 1965; Simon, 1997). For such one dimensional densities examples of different distributions with the same moment series have been known for more than a century (Heyde, 1963; Stieltjes, 1894).

It is not uncommonly argued in the cosmological literature that the relation between the moments, the characteristic functional and the density function can be inverted, suggesting that the density is always uniquely set by the N-point moments (Fry, 1985; Mo et al., 2010, e.g.). It is important to keep in mind that this holds only when the characteristic functional can be written as a convergent power series in the moments in a region around $J = 0$. As already mentioned this is not always true, in which case the characteristic functional cannot be expressed in terms of N-point moments. However, it is true that the mapping between the characteristic functional and the density is one to one. The indeterminacy of the moment problem was touched upon in a cosmological context in Coles and Jones (1991), though it did not attract much attention in the cosmological literature since then.

Obviously, an indeterminate moment problem is relevant for our purposes as it implies that the entire N-point function hierarchy contains less information than the density itself. In that case, the entire N-point hierarchy is an inefficient set of observables. Namely, it is impossible to reconstruct uniquely the density from the hierarchy. The implications for cosmological parameter inference are discussed in several chapters of this thesis, notably in chapter 6, where to the best of our knowledge first explicit examples of densities of any dimensionality with identical N-point moments at all orders are presented.

The Gaussian field is an example of a density that can be uniquely recovered from the N-point moment hierarchy. On the other hand, the lognormal field, first introduced later in 1.4 is an example where this is not possible.

1.3.3 N-point functions from discrete populations, poisson samples

In galaxy or weak lensing surveys, the fields that are observed are rather discrete than continuous. Namely positions of galaxies are recorded, and additional information such as the distortion of galaxy images can be effectively measured only on these positions. Discreteness adds some complexity. We are not able to predict galaxy positions in the sky, and this discrete field may not trace in an obvious manner the underlying, interesting field, typically the dark matter density field or its projection. Rather, they are only tracers

that can be biased in several ways, and the measured N-point functions need not always be representative of that of the underlying field.

There is no unique manner to create a point process from a continuous random field, but for density fields the infinitesimal Poisson model is rather natural and gives a direct interpretation of the point functions. Within this prescription, one divides the total volume in infinitesimal cells and simply set the probability for a point in a cell to be proportional to the value of the continuous field ϕ at that point, and this independently from cells to cells. In that case, given a total number of N objects in a total volume V, the probability density to find these at $x_1 \cdots, x_N$ conditional on ϕ is by definition

$$p(x_1, \cdots, x_N|\phi) = \phi(x_1) \cdots \phi(x_N) \frac{1}{\left(\bar{\phi}V\right)^N}. \tag{1.25}$$

Note that the normalisation $(\bar{\phi}V)^N$ must require ergodicity, in order to be able to identify $\int d^n x_i \, \phi(x_i)$ with $\bar{\phi}V$. This condition implies that for any k the following link between N-point functions must hold

$$\frac{1}{\bar{\phi}V} \int_V d^n x \, \langle \phi(x)\phi(x_1) \cdots \phi(x_k) \rangle = \langle \phi(x_1) \cdots \phi(x_k) \rangle \,, \tag{1.26}$$

which is a very non trivial condition on a density function, requiring in fact the zeroth mode of the field, $\int_V d^n x \, \phi(x) = \tilde{\phi}(0)$ to be actually no random variable but a usual number. For a Gaussian field, this is equivalent to require the condition on the spectrum $P(0) = 0$ and thus presents no difficulty for any volume V. On the other hand, as we will discuss in chapter 4, a lognormal field never has this property fulfilled exactly in a finite volume, since its power at zero must be strictly positive.

We find the probability density $p(x_1, \cdots, x_N)$, unconditional to ϕ, to find these objects at these positions by marginalising over the unseen underlying field ϕ,

$$\begin{aligned} p(x_1, \cdots, x_N) &= \int \mathcal{D}\phi \, p[\phi] \, p(x_1, \cdots, x_N|\phi) \\ &= \frac{1}{\left(\bar{\phi}V\right)^N} \langle \phi(x_1) \cdots \phi(x_N) \rangle \,. \end{aligned} \tag{1.27}$$

From (1.27) we find a direct interpretation of the two-point function. From the rule of probability theory we find that the probability density of observing an object at x_2 given that there is one at x_1 is given by

$$p(x_2|x_1) d^n x_2 = \frac{1}{\bar{\phi}^2 V} \langle \phi(x_1)\phi(x_2) \rangle \, d^n x_2 = \left(1 + \xi\left(x_2 - x_1\right)\right) \frac{d^n x_2}{V}. \tag{1.28}$$

Thus for such processes the connected two-point function describes directly the clustering of the points, by enhancing or reducing this conditional probability to find particles separated by some distance (Bernardeau et al., 2002; Peebles, 1980, e.g.).

1.4 Gaussian and non-Gaussian matter density fields

In cosmology, the Gaussian field (1.22) is fundamental. It is used routinely in order to describe the statistics of the small fluctuations present in the early Universe that we observe in the CMB radiation, or more importantly for us that of the density we observe on the largest scales. While it is possible to treat this working hypothesis as an ad-hoc assumption adopted for convenience or lack of a better prescription, it has now some theoretical support as well, in that the simplest model of inflation predict initial conditions that are extremely close to Gaussian (Liddle and Lyth, 2000). It is fortunately possible to see where this comes about without entering any details : in such models a nearly free scalar field is responsible for the rapid expansion of the Universe, and its fluctuations give rise to the primordial deviations from homogeneity. Since the action of a free field is quadratic in the field, and since the action plays the same role in a quantum field theory as $\ln p[\phi]$ in a statistical field theory, we see that it corresponds to a Gaussian field.

Nevertheless, it should be noted that in the particular case of the matter density or fluctuation field relevant for galaxy surveys, the assumption of Gaussianity is in fact flawed from the very beginning. This is because the matter density is positive, while the Gaussian assumption assigns non-zero probability density to negative values. As long as the variance is small, this is however not an essential shortcoming of the model.

Of course, there are many situations where non-Gaussian statistics play a major role, even in the noise-free fields. For instance, signatures from non-Gaussianities in the primordial fluctuations can be used to try and constrain more sophisticated inflationary models. This is often dubbed as primordial non-Gaussianity. In this thesis we are going to deal with the non-Gaussianity sourced from the nonlinear evolution of the density field. We already mentioned that Gaussian statistics cannot provide a perfect description of a density field, and this is even more true as nonlinear evolution take place. This is illustrated in figure 1.2, from A. Pillepich (Pillepich et al., 2010). The four panels on the left show the evolution of the matter density field in a N-body simulation, from redshift 50 to redshift 0 downwards. The inbox inside each panel shows the one-point probability density $p(1+\delta)$ of the matter fluctuation field as the dark line, together with that of the previous panel in grey. As the fluctuations grows, one observes the field to develop large tails in the overdense regions, and a cutoff in the underdense regions. The right panels are the same simulations where an amount of primordial non Gaussianity of the local type roughly 10 times larger than current observational constraints (Komatsu et al., 2011) was added, represented as the blue line in the inboxes. It is obvious that the non Gaussianity induced by the formation of structures is both very strong and completely dominant over the primordial non-Gaussianity in the late Universe.

More to the point in the case of the density field is the assumption of a lognormal field (Coles and Jones, 1991), where essentially the logarithm $A = \ln(1+\delta)$ is set to be a Gaussian field. Since $\ln(1+\delta)$ is very close to δ for small fluctuations, these two fields are indistinguishable for any practical purposes as long as the variance is small. The lognormal is however always positive definite, correcting for the defect of the Gaussian prescription.

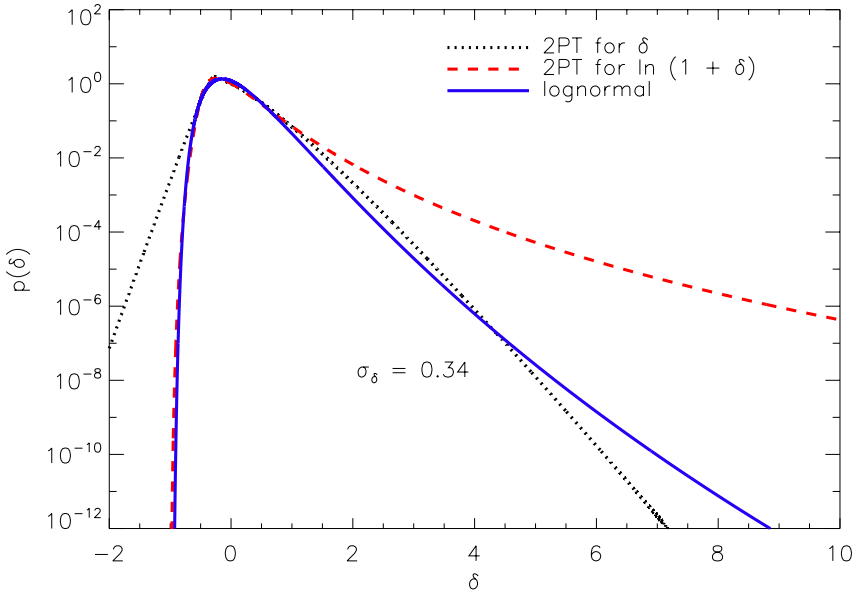

Fig. 1.1 — The shape of the one-point distribution function of the fluctuation field according to second order perturbation theory for δ, dotted, and for $\ln(1+\delta)$, dashed, calculated with methods exposed in Taylor and Watts (2000). The solid line is the lognormal distribution.

In the nonlinear regime, it also shows large tails and a cutoff in the underdense regions, reproducing the qualitative features the one-point distribution of the fluctuation field remarkably well given the simplicity of the prescription.

Let us illustrate these aspects in figure 1.1 with the help of perturbation theory. As the dotted line is shown the prediction of second order perturbation theory for the distribution function, calculated with the same methods as Taylor and Watts (2000), assuming Gaussian initial conditions for δ, at a variance of $\sigma_\delta = 0.34$. We see that the assumption of Gaussian initial conditions gives rise to a nonsensical non zero probability density for negative matter density. On the other hand, as expected, it predicts a large tail in the overdense regions. The dashed line shows the probability density obtained from the same type of perturbative calculations for A, and the solid line the lognormal distribution. It is rather remarkable how both the tails and the underdense regions of the lognormal are reproduced with these two perturbative calculations. The agreement between the lognormal prescription and the matter fluctuations measured in the N-body simulations is also seen to be very good, see Taylor and Watts (2000).

We can see from the apparition of a long tail and a sharp cutoff in the density fluctuations that the statistics of the nonlinear regime are not going to obey the same rules as that of the linear Gaussian field. For instance, it is worth mentioning at this point that the traditional observable the spectrum of the field, that contains the entire information in the field in the Gaussian regime, seems to lose most of its advantages leaving the nonlinear

regime. In particular, heavy correlations appear between the Fourier modes, and little information can apparently be extracted from the spectrum on these scales (Lee and Pen, 2008; Neyrinck et al., 2006; Rimes and Hamilton, 2005). The study of the information within the lognormal field will form a large part of this thesis, chapter 6.

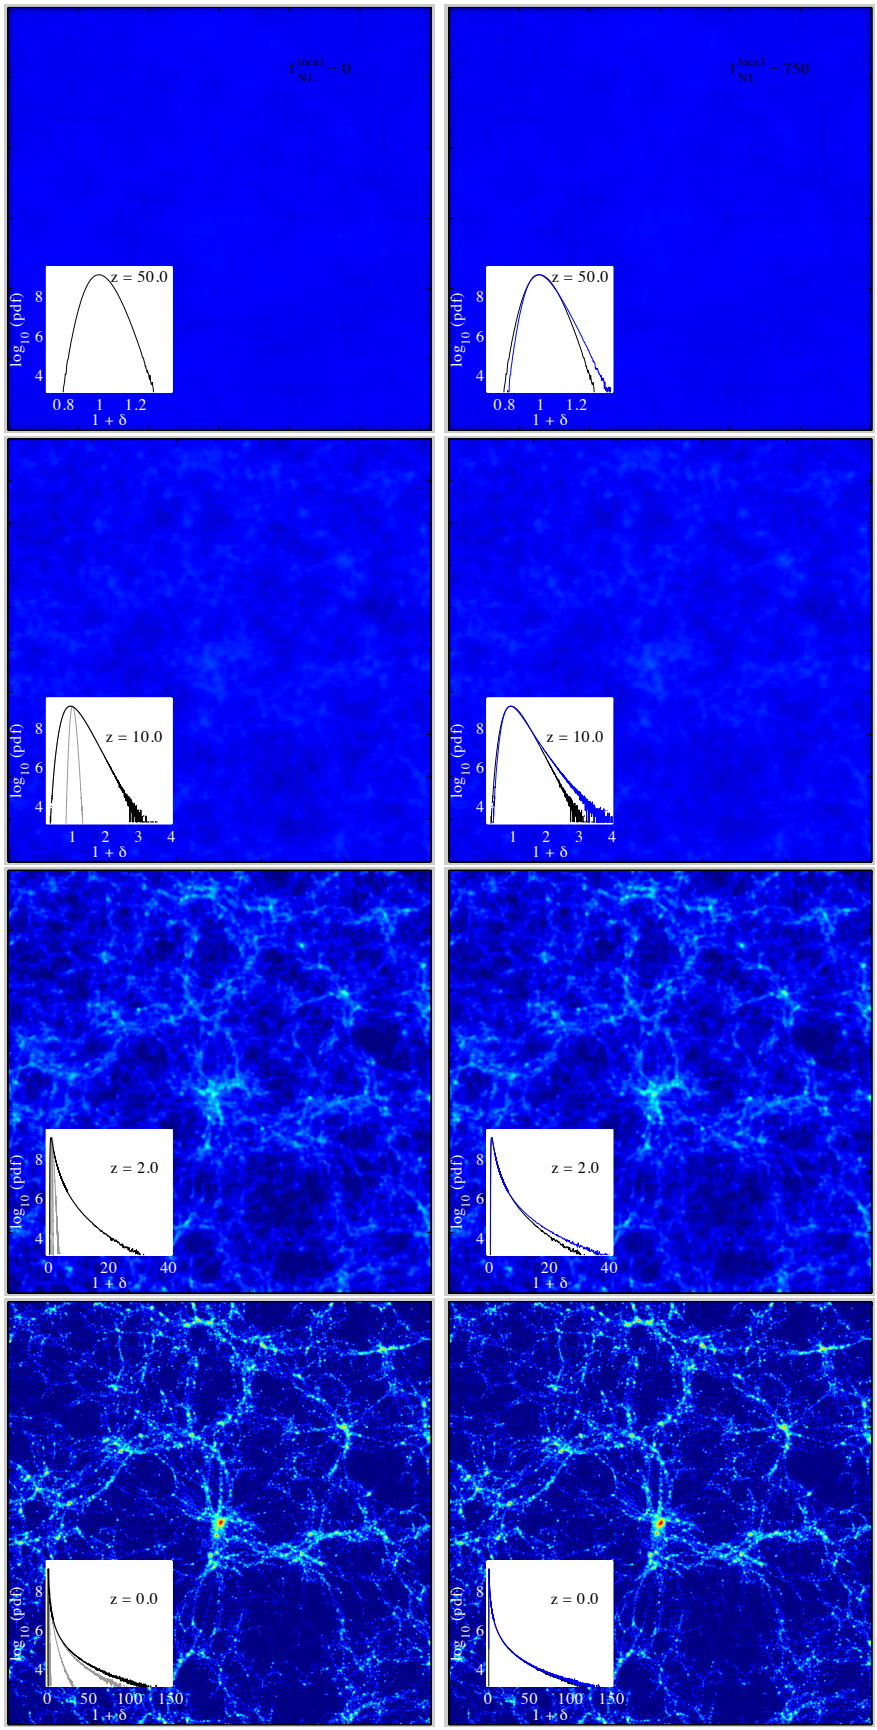

Fig. 1.2 — Evolution of the matter density field distribution with cosmic time in a N-body simulation. See text for more details. Image courtesy of A. Pillepich (Pillepich et al., 2010).

1.5 Structure of the thesis

The thesis is built out of two parts. The first part, that includes the first and second chapter, describes and builds upon the mathematical tools that are used in later chapters. The second part, chapter 3 to chapter 7, contains the cosmological research properly speaking. Each chapter begins with a more detailed description of its content, as well as the references to our corresponding publications when appropriate.

In chapter 2, we discuss two measures of information, Fisher's information matrix and Shannon's information entropy, and the duality between them. We comment extensively on the information inequality, fundamental for much of this thesis. We discuss the information content of maximal entropy distributions, and identify these distributions as those for which the information inequality is an equality.

In chapter 3 we deal with the information content of N-point moments for a given density function. We present how to decompose the Fisher information matrix in uncorrelated components, associated to the N-point moments of a given order, with the help of orthogonal polynomials. The properties of this expansion are discussed. In particular, it is shown that the hierarchy of N-point moments does not necessarily carry the entire Fisher information content of the distribution for indeterminate moment problems, while the entire information is recovered for determinate moment problems. We also present several models for which we could obtain this expansion explicitly at all orders.

Chapter 4 uses the tools introduced in chapter 2 in the context of weak lensing. We show in this chapter how the information from different probes of the weak lensing convergence field, the magnification, shears and flexion fields, do combine in a very simple way. We then evaluate their information content using current values for the dispersion parameters and discuss the benefits of their combination.

Chapter 5 is a little note on the use of Gaussian distributions for the statistics of estimators of second order statistics in cosmology. We show that we can use these information-theoretic concepts to clarify some issues in the literature and the signification of the parameter dependence of the covariance matrices for two-point correlation functions or power spectra.

Chapter 6 discusses the information content of N-point moments in the lognormal density field. It is the largest chapter in this thesis, and contains our main results as well. After discussing some fundamental limitations of the lognormal field to describe the matter density field of the ΛCDM universe, we present families of different fields all having the same hierarchy of N-point moments than the lognormal, and discuss the implications for parameter inference. The expansion of the Fisher information matrix introduced in chapter 2 is then performed exactly at all orders in two simplified but tractable situations. This then allows us to make successful connections with N-body simulation results on the extraction of power spectra

Finally, in chapter 7 we evaluate the information content of the moment hierarchy of the one-point distribution of the weak lensing convergence field, demonstrating that the nonlinearities generically lead to distributions that are very poorly described by their moments. On the other hand, it is shown that simple mappings are able to correct for this deficiency.

Part I

Quantifying information

2 Shannon entropy and Fisher information

The primary aim of this first chapter is to introduce in a rather detailed and comprehensive manner the tools that form the building blocks of this thesis, which are Fisher's matrix valued measure of information, as well as Shannon's measure of entropy. As such, unlike the subsequent chapters, it does not contain exclusively original material. Namely, the sections 2.1 and 2.2, while important parts of our publication Carron et al. (2011) can be considered to some extent a review of our perspective on the properties of Fisher information and Shannon entropy that are then built upon in the later parts of this thesis.

This chapter is built as follows :

In section 2.1, we introduce and discuss the main properties of the Fisher information, with an emphasis on its information theoretic properties. Essential to most of this thesis is the information inequality, equation (2.21), and its consequences. In section 2.2, we discuss maximal entropy distributions associated to a prescribed set of observables, and that these distributions are precisely those for which the information inequality is an equality. We find with equation (2.44) a measure of information content that depends only on the constraints put on the data and the physical model, written in terms of the curvature of Shannon's entropy surface. We recover the Fisher information matrices for Gaussian fields of common use in cosmology as the special case of fields with prescribed two-point functions.

The text in these two sections is based to a large extent on the first part of Carron et al. (2011), with the exception of appendix 2.3.1.

2.1 Fisher information and the information inequality

We first review here a few simple points of interest that justify the interpretation of the Fisher matrix as a measure of the information content of an experiment. Let us begin by considering the case of a single measurement X, with different possible outcomes, or realisations, x, and our model has a single parameter α. We also assume that we have knowledge, prior to the given experiment, of the probability density function $p_X(x, \alpha)$, which depends on our parameter α, that gives the probability of observing particular realisations for each value of the model parameter. The Fisher information, F, in X on α, is a non-negative scalar in this one parameter case. It is defined in a fully general way as a sum over all realisations of the data (Fisher, 1925):

$$I_X(\alpha) = \left\langle \left(\frac{\partial \ln p_X(x, \alpha)}{\partial \alpha} \right)^2 \right\rangle . \tag{2.1}$$

Three simple but important properties of Fisher information are worth highlighting at this point.

- The first is that $I_X(\alpha)$ is positive definite, and it vanishes if and only if the parameter α does not impact the data, i.e. if the derivative of $p_X(x, \alpha)$ with respect to α is zero for every realisation x.

- The second point is that it is invariant to invertible manipulations of the observed data. This can be seen by considering an invertible change of variable $y = f(x)$, which, due to the rules of probability theory can be expressed as

  $$p_Y(y, \alpha) = p_X(x, \alpha) \left| \frac{dx}{dy} \right| . \tag{2.2}$$

  Thus

  $$\frac{\partial \ln p_Y(y, \alpha)}{\partial \alpha} = \frac{\partial \ln p_X(x, \alpha)}{\partial \alpha}, \tag{2.3}$$

  leading to the simple equivalence that $I_X(\alpha) = I_Y(\alpha)$. On the other hand, information may be lost when the transformation is not unique in both directions. For instance, if the data is combined to produce a new variable that could arise from different sets of data points. This is only the statement that manipulations of the data leads, at best, only to conservation of the information.

- The third point is that information from independent experiments add together. Indeed, if two experiments with data X and Y are independent, then the joint probability density factorises,

  $$p_{XY}(x, y) = p_X(x) p_Y(y), \tag{2.4}$$

  and it is easy to show that the joint information in the observations decouples,

  $$I_{XY}(\alpha) = I_X(\alpha) + I_Y(\alpha). \tag{2.5}$$

These three properties satisfy what we might intuitively expect from a mathematical implementation of an abstract concept such as information. However, we can ask the reverse question and try to find an alternative measure of information that may be better suited, for a particular purpose, than the Fisher information. In appendix 2.3.1, we discuss to what extent the information measure in (2.1) is in fact uniquely set by the these requirements above.

These properties are making the Fisher information a meaningful measure of information. This is independent of its interpretation as providing error bars on parameters. It further implies that once a physical model is specified with a given set of parameters, a given experiment has a definite information content that can only decrease with data processing.

2.1.1 The case of a single observable

To quantify the last point above, and in order to get an understanding of the structure of the information in a data set, we first discuss a simple situation, common in cosmology, where the extraction of the model parameter α from the data goes through the intermediate step of estimating a particular observable, D, from the data, x, with the help of which α will be inferred. A typical example could be, from the temperature map of the CMB (x), the measurement of the power spectra of the fluctuations (D), from which a cosmological parameter (α) is extracted. The observable D is measured from x with the help of an estimator, that we call $\hat{D}$, and that we will take as unbiased. This means that its mean value, as would be obtained for instance if many realizations of the data were available, converges to the actual value that we want to compare with the model prediction,

$$\left\langle \hat{D} \right\rangle = D(\alpha). \tag{2.6}$$

A measure for its deviations from sample to sample, or the uncertainty in the actual measurement, is then given by the variance of $\hat{D}$, defined as

$$\mathrm{Var}\left(\hat{D}\right) = \left\langle \hat{D}^2 \right\rangle - \left\langle \hat{D} \right\rangle^2. \tag{2.7}$$

In such a situation, a major role is played by the so-called Cramér-Rao inequality (Rao (1973)), that links the Fisher information content of the data to the variance of the estimator, stating that

$$\mathrm{Var}(\hat{D}) I_X(\alpha) \geq \left(\frac{\partial D(\alpha)}{\partial \alpha} \right)^2. \tag{2.8}$$

This equation holds for any such estimator $\hat{D}$ and any model parameter α. Two different interpretations of this equation are possible:

The first bounds the variance of $\hat{D}$ by the inverse of the Fisher information. To see this, we consider the special case of the model parameter α being D itself. Although we

are making in general a conceptual distinction between the observable D and the model parameter α, nothing requires us from doing so. Since α is now equal to D, the derivative on the right hand side becomes unity, and one obtains

$$\mathrm{Var}(\hat{D}) \geq \frac{1}{I_X(D)}. \tag{2.9}$$

The variance of any unbiased estimator $\hat{D}$ of D is therefore bounded by the inverse of the amount of information $I_X(D)$ the data possess on D. If $I_X(D)$ is known it gives a useful lower limit on the error bars that the analysis of the data can put on this observable. We emphasise at this point that this bound only holds in the case of unbiased estimators. There are cosmologically relevant situations where biased estimators can go beyond this level, and thus perform better according to the minimal squared error criterion than any unbiased one.

The second reading of the Cramér-Rao inequality, closer in spirit to the present thesis, is to look at how information is lost by constructing the observable D, and discarding the rest of the data set. For this, we rewrite trivially equation (2.8) as

$$I_X(\alpha) \geq \left(\frac{\partial D}{\partial \alpha}\right)^2 \frac{1}{\mathrm{Var}(\hat{D})}. \tag{2.10}$$

The expression on the right hand side is the ratio of the sensitivity of the observable to the model parameter $\left(\frac{\partial D}{\partial \alpha}\right)^2$, to the accuracy with which the observable can be extracted from the data, $\mathrm{Var}(\hat{D})$. One of the conceivable approaches in order to estimate the true value of the parameter α, is to perform a χ^2 fit to the measured value of D . It is simple to show that this ratio, evaluated at the best fit value, is in fact proportional to the expected value of the curvature of $\chi^2(\alpha)$ at this value. Since the curvature of the χ^2 surface describes how fast the value of the χ^2 is increasing when moving away from the best fit value, its inverse may be interpreted as an approximation to the error estimate that the analysis with the help of $\hat{D}$ will put on α.

Thus, equation (2.10) shows that by only considering D and not the full data set, we may have lost information on α, a loss given by the difference between the left and right hand side of that equation. While the latter may be interpreted as the information on α contained in the part of the data represented by D, we may have lost trace of any other source of information.

It should be noted that while we have just chosen to interpret the right hand side of (2.10) as the information in $\hat{D}$, this is a slight abuse of terminology. More rigorously, the Fisher information in $\hat{D}$ is not the right hand side of (2.10) but the Fisher information of its density function

$$I_{\hat{D}}(\alpha) = \left\langle \left(\frac{\partial \ln p_{\hat{D}}}{\partial \alpha}\right)^2 \right\rangle, \quad p_{\hat{D}}(\hat{D}) = \int dx\, p_X(x) \delta^D(\hat{D} - \hat{D}(x)), \tag{2.11}$$

but it will always be clear in this thesis from the context which one is meant. Anticipating the nomenclature of chapter 3, the right hand side of (2.10) is actually the information in

the mean of $\hat{D}$. From equation (2.10) we infer

$$I_{\hat{D}}(\alpha) \geq \left(\frac{\partial D}{\partial \alpha}\right)^2 \frac{1}{\mathrm{Var}(\hat{D})}. \tag{2.12}$$

2.1.2 The general case

These considerations on the Cramér-Rao bound can be easily generalised to the case of many parameters and many estimators of as many observables. Still dealing with a measurement X with outcomes x, we want to estimate a set of parameters

$$\boldsymbol{\theta} = (\alpha, \beta, \cdots) \tag{2.13}$$

with the help of some vector of observables,

$$\mathbf{D} = (D_1, \cdots, D_n) \tag{2.14}$$

that are extracted from x with the help of an array of unbiased estimators,

$$\hat{\mathbf{D}} = \left(\hat{D}_1, \cdots, \hat{D}_n\right), \quad \left\langle \hat{\mathbf{D}} \right\rangle = \mathbf{D} \tag{2.15}$$

In this multidimensional setting, all the three scalar quantities that played a role in our discussion in section 2.1.1, i.e. the variance of the estimator, the derivative of the observable with respect to the parameter, and the Fisher information, are now matrices.

The Fisher information F in X on the parameters $\boldsymbol{\theta}$ is defined as the square matrix

$$\left[F^X(\boldsymbol{\theta})\right]_{\alpha\beta} = \left\langle \frac{\partial \ln p_X}{\partial \alpha} \frac{\partial \ln p_X}{\partial \beta} \right\rangle. \tag{2.16}$$

While the diagonal elements $F^X_{\alpha\alpha}$ are the information scalars $I_X(\alpha)$ in equation (2.1), the off diagonal ones describe correlated information. The Fisher information matrix still carries the three properties we discussed in section 2.1.

The variance of the estimator in equation (2.7) now becomes the covariance matrix $\mathrm{cov}(\hat{\mathbf{D}})$ of the estimators $\hat{\mathbf{D}}$, defined as

$$\mathrm{cov}\left(\hat{\mathbf{D}}\right)_{ij} = \left\langle \hat{D}_i \hat{D}_j \right\rangle - D_i D_j. \tag{2.17}$$

Finally, the derivative of the observable with respect to the parameter, in the right hand side of (2.8), becomes a matrix Δ, in general rectangular, defined as

$$\Delta_{\alpha\, i} = \frac{\partial D_i}{\partial \alpha}, \tag{2.18}$$

where α runs over all elements of the set $\boldsymbol{\theta}$ of model parameters. Again, the Cramér-Rao inequality provides a useful link between these three matrices, and again there are two

approaches to that equation : first, as usually presented in the literature (Rao, 1973), in the form of a lower bound to the covariance matrix of the estimators,

$$\mathrm{cov}\left(\hat{\mathbf{D}}\right) \geq \Delta^T \left[F^X\left(\boldsymbol{\theta}\right)\right]^{-1} \Delta. \tag{2.19}$$

The inequality between two symmetric matrices $A \geq B$ having the meaning that the matrix $A - B$ is positive definite. [1]. If, as above, we consider the special case of identifying the parameters with the observables themselves, the matrix Δ is the identity matrix, and so we obtain that the covariance of the vector of the estimators is bounded by the inverse of the amount of Fisher information that there is on the observables in the data,

$$\mathrm{cov}(\hat{\mathbf{D}}) \geq \left[F^X(\mathbf{D})\right]^{-1}. \tag{2.20}$$

Second, we can turn this lower bound on the covariance to a lower bound on the amount of information in the data set as well. By rearranging equation (2.19), we obtain the multidimensional analogue of equation (2.10), the information inequality, which describes the loss of information that occurs when the data is reduced to a set of estimators,

$$F^X\left(\boldsymbol{\theta}\right) \geq \Delta \left[\mathrm{cov}\left(\hat{\mathbf{D}}\right)\right]^{-1} \Delta^T. \tag{2.21}$$

This information inequality is a central piece to much of this thesis. A proof can be found in the appendix. A maybe simpler proof follows also for instance from the discussion in the appendix of chapter 3.

Instead of giving a useful lower bound to the covariance of the estimator as in the Cramér-Rao inequality, equation (2.19), the information inequality makes clear how information is in general lost when reducing the data to any particular set of estimators. The right hand side may be seen, as before, as the expected curvature of a χ^2 fit to the estimates produced by the estimators $\hat{\mathbf{D}}$, when evaluated at the best fit value, with all correlations fully and consistently taken into account. Note that as before the right hand side of the information inequality is not the Fisher information content of the joint probability density function of the estimators, but only that of their means.

In section 2.2, we discuss how Jaynes' Maximal Entropy Principle allow us to understand the total information content of a data set, once a model is specified, in very similar terms.

2.1.3 Resolving the density function: Fisher information density

Due to its generality, the information inequality (2.21) is very powerful. We now have a deeper look at a special case that sheds some light on the definition of the Fisher infor-

[1] A symmetric matrix A is called positive definite when for any vector x holds that $x^T A x \geq 0$. Further we say for such matrices that A is larger than B, or $A > B$, whenever $A - B > 0$. A concrete implication for our purposes is e.g. that the diagonal entries of the left hand side of (2.19) or (2.20), which are the individual variances of each estimator $\hat{D}_i$, are greater than those of the right hand side. For many more properties of positive definite matrices, see for instance (Bhatia, 2007)

mation matrix (2.16), and that we will use in part II.

Assuming that the variable x is continuous and one dimensional, pick a set of points x_i with separation dx covering some range A of the variable, such that in the limit of a large number of points we can write

$$\sum_i p_X(x_i)dx \to \int_A dx\, p_X(x), \tag{2.22}$$

Generalisation to discrete variables or multidimensional cases will be obvious.
Define then a set of estimator $\hat{D}_i$ as follows

$$\hat{D}_i(x) := \begin{cases} 1, & x \in (x_i, x_i + dx) \\ 0, & x \notin (x_i, x_i + dx) \end{cases}. \tag{2.23}$$

These estimators simply build an histogram of the variable over A. In other words, our set of estimators are defined such the entire density function is resolved over A.

We want to evaluate the information inequality for this set of estimators. We have

$$\left\langle \hat{D}_i \right\rangle = p_X(x_i)dx, \tag{2.24}$$

with covariance matrix

$$\mathrm{cov}\left(\hat{\mathbf{D}}\right)_{ij} = \delta_{ij} p_X(x_i)dx - dx\, p_X(x_i)\, dx\, p_X(x_j). \tag{2.25}$$

Define as p_X^- the probability that a realisation of the variable does not belong to A :

$$p_X^- := \int_{\mathbb{R}/A} dx\, p_X(x). \tag{2.26}$$

It is then easily seen that the inverse covariance matrix is

$$\left[\mathrm{cov}\left(\hat{\mathbf{D}}\right)^{-1}\right]_{ij} = \frac{\delta_{ij}}{p_X(x_i)dx} + \frac{1}{p_X^-}. \tag{2.27}$$

It follows that the right hand side of the information inequality becomes

$$\sum_{i,j} (dx)^2 \frac{\partial p_X(x_i)}{\partial \alpha} \left[\mathrm{cov}\left(\hat{\mathbf{D}}\right)^{-1}\right]_{ij} \frac{\partial p_X(x_j)}{\partial \beta} = \int_A dx\, \frac{1}{p_X(x)} \frac{\partial p_X(x)}{\partial \alpha} \frac{\partial p_X(x)}{\partial \beta} + \frac{1}{p_X^-} \frac{\partial p_X^-}{\partial \alpha} \frac{\partial p_X^-}{\partial \beta}. \tag{2.28}$$

This is nothing else than

$$\int_A dx\, p_X(x) \frac{\partial \ln p_X(x)}{\partial \alpha} \frac{\partial \ln p_X(x)}{\partial \beta} + p_X^- \frac{\partial \ln p_X^-}{\partial \alpha} \frac{\partial \ln p_X^-}{\partial \beta}. \tag{2.29}$$

If A covers the full range of the variable, then the first term is precisely the total Fisher information matrix, and the second vanishes, since $p_X^- = 0$. It is clear from (2.29) that we can interpret

$$p_X(x) \frac{\partial \ln p_X(x)}{\partial \alpha} \frac{\partial \ln p_X(x)}{\partial \beta} = \frac{1}{p_X(x)} \frac{\partial p_X(x)}{\partial \alpha} \frac{\partial p_X(x)}{\partial \beta} \tag{2.30}$$

as a Fisher information density, representing information from observations of the variables around x. The additional term in (2.29) involving p_X^- originates from the fact the the density is normalised to unity : observations of the density over the range A provides some information on the density in the complement to A,

$$p_X^- = \int_{\mathbb{R}/A} dx\, p_X(x) = 1 - \int_A dx\, p_X(x). \tag{2.31}$$

However, it is not possible to resolve the individual contributions of $p_X(x)$ to p_X^- for each x on the complement of A, and thus the derivatives act in this case outside of the integrals, unlike the first term in (2.29).

2.2 Jaynes Maximal Entropy Principle

In cosmology, the knowledge of the probability distribution of the data as function of the parameters, $p_X(x, \boldsymbol{\theta})$, which is compulsory in order to evaluate its Fisher information content, is usually very limited. In a galaxy survey, a data outcome x would be typically the full set of angular positions of the galaxies, together with some redshift estimation if available, to which we may add any other kind of information, such as luminosities, shapes, etc. Our ignorance of both initial conditions and of many relevant physical processes does not allow us to predict either galaxy positions in the sky, or all interconnections with all this additional information. Our predictions of the shape of p_X is thus limited to some statistical properties, that are sensitive to the model parameters $\boldsymbol{\theta}$, such as the mean density over some large volume, or certain types of correlation functions.

In fact, even if it were possible to devise some procedure in order to get the exact form of p_X, it may eventually turn out to be useless, or even undesirable, to do so. The incredibly large number of degrees of freedom of such a function is very likely to overwhelm the analyst with a mass of irrelevant details, which may have no relevant significance on their own, or improve the analysis in any meaningful way.

These arguments call for a kind a thermodynamical approach, which would try and capture those aspects of the data which are relevant to our purposes, reducing the number of degrees of freedom in a drastic way. Such an approach already exists in the field of probability theory (Jaynes, 1957). It is based on Shannon's concept of entropy of a probability distribution (Shannon, 1948) and did shed new light on the connection between probability theory and statistical mechanics.

As we have just argued, our predictive knowledge of $p_X(x, \boldsymbol{\theta})$ is limited to some statistical properties. Let us formalise this mathematically, in a similar way as in section 2.1.2. Astrophysical theory gives us a set of constraints on the shape of p_X, in the form of averages of some functions o_i,

$$O_i(\boldsymbol{\theta}) = \langle o_i(x) \rangle (\boldsymbol{\theta}), \quad i = 1, \cdots, n. \tag{2.32}$$

where p_X enters through the angle brackets. As an example, suppose the data outcome x is a map of the matter density field as a function of position. In this case, one of these constraints O_i could be the mean of the field or its power spectrum, as given by some cosmological model.

The role of this array $\mathbf{O} = (O_1, \cdots, O_n)$ is to represent faithfully the physical understanding we have of p_X, according to the model, as a function of the model parameters $\boldsymbol{\theta}$. In the ideal case, some way can be devised to extract each one of these quantities O_i from the data and to confront them to theory. The set of observables $\mathbf{D}$, that we used in section 2.1.2, would be a subset of these predictions $\mathbf{O}$, and we henceforth refer to $\mathbf{O}$ as the 'constraints'.

Although p_X must satisfy the constraints (2.32), there may still be a very large number of different distributions compatible with these. However, a very special status among these distributions has the one which maximises the value of Shannon's entropy[2], defined as

$$S = -\int dx\, p_X(x, \boldsymbol{\theta}) \ln p_X(x, \boldsymbol{\theta}). \tag{2.33}$$

First introduced by Shannon (Shannon, 1948) as a measure of the uncertainty in a distribution on the actual outcome, Shannon's entropy is now the cornerstone of information theory. Jaynes' Maximal Entropy Principle states that the p_X for which this measure S is maximal is the one that best deals with our insufficient knowledge of the distribution, and should be therefore preferred. We refer the reader to Jaynes' work (Jaynes, 1983; Jaynes and Bretthorst, 2003) and to Caticha (2008) for detailed discussions of the role of entropy in probability theory and for the conceptual basis of maximal entropy methods. Astronomical applications related to some extent to Jaynes's ideas include image reconstruction from noisy data, (see e.g. Maisinger et al. (2004); Skilling and Bryan (1984); Starck and Pantin (1996) and references therein) , mass profiles reconstruction from shear estimates (Bridle et al., 1998; Marshall et al., 2002), as well as model comparison when very few data is available (Zunckel and Trotta, 2007). We will see that for our purposes as well it provides us a powerful tool, and that the Maximal Entropy Principle is the ideal complement to Fisher information, fitting very well within our discussions in section 2.1 on the information inequality.

Intuitively, the entropy S of p_X tells us how sharply constrained the possible outcomes x are, and Jaynes' Maximal Entropy Principle selects the p_X which is as wide as possible, but at the same time consistent with the constraints (2.32) that we put on it. The actual maximal value attained by the entropy S, among all the possible distributions which satisfy (2.32), is a function of the constraints $\mathbf{O}$, which we denote by

$$S(O_1, \cdots, O_n). \tag{2.34}$$

[2]Formally, for continuous distributions the reference to another distribution is needed to render S invariant with respect to invertible transformations, leading to the concept of the entropy of p_X relative to another distribution q_X, $S = \int dx\, p_X(x) \ln \frac{p_X(x)}{q_X(x)}$, also called Kullback-Leibler divergence. The quantity defined in the text is more precisely the entropy of $p_X(x)$ relative to a uniform probability density function. For an recent account on this, close in spirit to this work, see Caticha (2008).

Of course it is a function of the model parameters $\boldsymbol{\theta}$ as well, since they enter the constraints. As we will see, the shape of that surface as a function of $\mathbf{O}$, and thus implicitly as a function of $\boldsymbol{\theta}$, is the key point in understanding the Fisher information content of the data. In the following, in order to keep the notation simple, we will omit the dependency on $\boldsymbol{\theta}$ of most of our expressions, though it will always be implicit.

The problem of finding the distribution p_X that maximises the entropy (2.33), while satisfying the set of constraints (2.32), is an optimization exercise. We can quote the end result (Jaynes, 1983, chap. 11),(Caticha, 2008, chap. 4):
The probability density function p_X, when it exists, has the following exponential form,

$$p_X(x) = \frac{1}{Z} \exp\left(-\sum_{i=1}^{n} \lambda_i o_i(x) \right), \tag{2.35}$$

in which to each constraint O_i is associated a conjugate quantity λ_i, that arises formally as a Lagrange multiplier in this optimization problem with constraints. The conjugate variables λ's are also called 'potentials', terminology that we will adopt in the following. We will see below in equation (2.39) that the potentials have a clear interpretation, in the sense that the each potential λ_i quantifies how sensitive is the entropy function S in (2.34) to its associated constraint O_i. The quantity Z, that plays the role of the normalisation factor, is called the partition function. Since equation (2.35) must integrate to unity, the explicit form of the partition function is

$$Z(\lambda_1, \cdots, \lambda_n) = \int dx \, \exp\left(-\sum_{i=1}^{n} \lambda_i o_i(x) \right). \tag{2.36}$$

The actual values of the potentials are set by the constraints (2.32). They reduce namely, in terms of the partition function, to a system of equations to solve for the potentials,

$$O_i = -\frac{\partial}{\partial \lambda_i} \ln Z, \quad i = 1, \cdots, n. \tag{2.37}$$

The partition function Z is closely related to the entropy S of p_X. It is simple to show that the following relation holds,

$$S = \ln Z + \sum_{i=1}^{n} \lambda_i O_i, \tag{2.38}$$

and the values of the potentials can be explicitly written as function of the entropy, in a relation mirroring equation (2.37),

$$\lambda_i = \frac{\partial S}{\partial O_i}, \quad i = 1, \cdots, n \tag{2.39}$$

Given the nomenclature, it is of no surprise that a deep analogy between this formalism and statistical physics does exist. Just as the entropy, or partition function, of a physical system determines the physics of the system, the statistical properties of these maximal

entropy distributions follow from the functional form of the Shannon entropy or its partition function as a function of the constraints. For instance, the covariance matrix of the constraints is given by

$$\langle (o_i(x) - O_i)\,(o_j(x) - O_j)\rangle = \frac{\partial^2 \ln Z}{\partial\lambda_i \partial\lambda_j} \tag{2.40}$$

In statistical physics the constraints can be the mean energy, the volume or the mean particle number, with potentials being the temperature, the pressure and the chemical potential. We refer to Jaynes (1957) for the connection to the physical concept of entropy in thermodynamics and statistical physics.

2.2.1 Information in maximal entropy distributions

With our choice of probabilities p_X given by equation (2.35), the amount of Fisher information on the parameters $\boldsymbol{\theta} = (\alpha, \beta, \cdots)$ of the model can be evaluated in a straightforward way. The dependence on the model goes through the constraints, or, equivalently, through their associated potentials. It holds therefore that

$$\begin{aligned} \frac{\partial \ln p_X(x)}{\partial\alpha} &= -\frac{\partial \ln Z}{\partial\alpha} - \sum_{i=1}^{n} \frac{\partial\lambda_i}{\partial\alpha} o_i(x) \\ &= \sum_{i=1}^{n} \frac{\partial\lambda_i}{\partial\alpha} \left[O_i - o_i(x)\right], \end{aligned} \tag{2.41}$$

where the second line follows from the first after application of the chain rule and equation (2.37). Using the covariance matrix of the constraints given in (2.40), the Fisher information matrix, defined in (2.16), can then be written as a double sum over the potentials,

$$F^X_{\alpha\beta} = \sum_{i,j=1}^{n} \frac{\partial\lambda_i}{\partial\alpha} \frac{\partial^2 \ln Z}{\partial\lambda_i \partial\lambda_j} \frac{\partial\lambda_j}{\partial\beta}. \tag{2.42}$$

There are several ways to rewrite this expression as a function of the constraints and/or their potentials. First, it can be written as a single sum by using equation (2.37) as

$$F^X_{\alpha\beta} = -\sum_{i=1}^{n} \frac{\partial\lambda_i}{\partial\alpha} \frac{\partial O_i}{\partial\beta}. \tag{2.43}$$

Alternatively, since we will be more interested in using the constraints as the main variables, and not the potentials, we can show, using equation (2.39), that it also takes the form [3]

$$F^X_{\alpha\beta} = -\sum_{i,j=1}^{n} \frac{\partial O_i}{\partial\alpha} \frac{\partial^2 S}{\partial O_i O_j} \frac{\partial O_j}{\partial\beta}. \tag{2.44}$$

[3] We note that this result is valid only for maximal entropy distributions and is not equivalent to the second derivative of the entropy with respect to the parameters themselves. However it is formally identical to the corresponding expression for the information content of distributions within the exponential family (Jennrich and Moore, 1975), or (van den Bos, 2007, chapter 4), once the curvature of the entropy surface is identified with the generalized inverse of the covariance matrix.

We will use both of these last expressions in chapter 4 of this thesis.

Equation (2.44) presents the total amount of information on the model parameters $\boldsymbol{\theta}$ in the data X, when the model predicts the set of constraints O_i. The amount of information is in the form of a sum of the information contained in each constraint, with correlations taken into account, as in the right hand side in equation (2.21). In particular, it is a property of the maximal entropy distributions, that if the constraints O_i are not redundant, then it follows that the curvature matrix of the entropy surface $-\partial^2 S$ is invertible and is the inverse of the covariance matrix $\partial^2 \ln Z$ between the observables. To see this explicitly, consider the derivative of equation (2.37) with respect to the potentials,

$$-\frac{\partial O_i}{\partial \lambda_j} = \frac{\partial^2 \ln Z}{\partial \lambda_i \partial \lambda_j}. \tag{2.45}$$

The inverse of the matrix on the left hand side, if it can be inverted, is $-\frac{\partial \lambda_i}{\partial O_j}$, which can be obtained taking the derivative of equation (2.39), with the result

$$-\frac{\partial \lambda_i}{\partial O_j} = -\frac{\partial^2 S}{\partial O_i \partial O_j}. \tag{2.46}$$

We have thus obtained in equation (2.44), combining Jaynes' Maximal Entropy Principle together with Fisher's information, the exact expression of the information inequality (2.21) for our full set of constraints, but with an equality sign.

We see that the choice of maximal entropy probabilities is fair, in the sense that all the Fisher information comes from what was forced upon the probability density function, i.e. the constraints. No additional Fisher information is added when these probabilities are chosen. In fact, as shown in the appendix this requirement alone is enough to single out the maximal entropy distributions, as being precisely those for which the information inequality is an equality. This can be understood in terms of sufficient statistics and goes back to Pitman and Wishart (1936) and Kopman (1936). For a discussion in the language of the exponential family of distribution see Zografos and Ferentinos (1994).

In the special case that the model parameters are the constraints themselves, we have

$$F^X_{O_i O_j} = -\frac{\partial^2 S}{\partial O_i O_j} = -\frac{\partial \lambda_i}{\partial O_j}, \tag{2.47}$$

which means that the Fisher information on the model predictions contained in the expected future data is directly given by the sensitivity of their corresponding potential. Also, the application of the Cramér-Rao inequality, in the form given in equation (2.20), to any set of unbiased estimators of $\mathbf{O}$, shows that the best joint, unbiased, reconstruction of $\mathbf{O}$ is given by the inverse curvature of the entropy surface $-\partial^2 S$, which is, as we have shown, $\partial^2 \ln Z$.

We emphasise at this point that although the amount of information is seen to be identical to the Fisher information in a Gaussian distribution of the observables with the above correlations, nowhere in our approach do we assume Gaussian properties. The distribution of the constraints $o_i(x)$ themselves is set by the maximal entropy distribution of the data.

2.2.2 Redundant observables

We have just seen that in the case of independent constraints, the entropy of p_X provides through equation (2.44) both the joint information content of the data, as well as the inverse covariance matrix between the observables. However, if the constraints put on the distribution are redundant, the covariance matrix is not invertible, and the curvature of the entropy surface cannot be inverted either. We show however that in these cases, our equations for the Fisher information content (2.42, 2.43, 2.44) are still fully consistent, dealing automatically with redundant information to provide the correct answer.

An example of redundant information occurs trivially if one of the functions $o_i(x)$ can be written in terms of the others. For instance, for galaxy survey data, the specification of the galaxy power spectrum as an constraint, together with the mean number of galaxy pairs as function of distance, and/or the two-points correlation function, which are three equivalent descriptions of the same statistical property of the data. Although the number of observables $\mathbf{O}$, and thus the number of potentials, describing the maximal entropy distribution greatly increases by doing so, it is clear that we should expect the Fisher matrix to be unchanged, by adding such superfluous pieces of information. A small calculation shows that the potentials adjust themselves so that it is actually the case, meaning that this type of redundant information is automatically discarded within this approach. Therefore, we need not worry about the independency of the constraints when evaluating the information content of the data, which will prove convenient in some cases.

There is another, more relevant type of redundant information, that allow us to understand better the role of the potentials. Consider that we have some set of constraints $\{O_i\}_{i=1}^n$, and that we obtain the corresponding p_X that maximises the entropy. This p_X could then be used to predict the value O_{n+1} of the average some other function $o_{n+1}(x)$, that is not contained in our set of predictions,

$$\langle o_{n+1}(x) \rangle =: O_{n+1}. \tag{2.48}$$

For instance, the maximal entropy distribution built with constraints on the first n moments of p_X, will predict some particular value for the $n+1$-th moment, O_{n+1}, that the model was unable to predict by itself.
Suppose now some new theoretical work provides the shape of O_{n+1} as a function of the model parameters. This new constraint can thus now be added to the previous set, and a new, updated p_X is obtained by maximising the entropy. There are two possibilities at this point :

- It may occur that the value of O_{n+1} as provided by the model is identical to the prediction by the maximal entropy distribution that was built without that constraint. Since the new constraint was automatically satisfied, the maximal entropy distribution satisfying the full set of $n+1$ constraints must be equal to the one satisfying the original set. From the equality of the two distributions, which are both of the form (2.35), it follows that the additional constraint must have vanishing associated

potential,

$$\lambda_{n+1} = 0, \tag{2.49}$$

while the other potentials are pairwise identical. It follows immediately that the total information, as seen from equation (2.43) is unaffected, and no information on the model parameters was gained by this additional prediction. A cosmological example would be to enforce on the distribution of some field, together with the two-points correlation function, fully disconnected higher order correlation functions. It is well known that the maximal entropy distribution with constrained two-points correlation function has a Gaussian shape, and that Gaussian distributions have disconnected points function at any order. No information is thus provided by these field moments of higher order in this case.
This argument shows that, for a given set of original constraints and associated maximal entropy distribution, any function $f(x)$, which was not contained in this set, with average F, can be seen as being set to zero potential. Such F's therefore do not contribute to the information.

- More interesting is, of course, the case where this additional constraint differs from the predictions obtained from the original set $\{O_i\}_{i=1}^n$. Suppose that there is a mismatch δO_{n+1} between the predictions of the maximal entropy distribution and the model. In this case, when updating p_X to include this constraint, the potentials are changed by this new information, a change given to first order by

$$\delta\lambda_i = \frac{\partial^2 S}{\partial O_i \partial O_{n+1}} \delta O_{n+1}, \quad i = 1, \cdots, n+1, \tag{2.50}$$

and the amount of Fisher information changes accordingly. It is interesting to note that the entropy itself is invariant at this order. From equation (2.39) we have namely

$$\delta S = \sum_{i=1}^{n+1} \lambda_i \, \delta O_i = \lambda_{n+1} \delta O_{n+1} = 0, \tag{2.51}$$

since the new constraint was originally at zero potential. The entropy is, therefore, stationary not only with respect to changes in the probability distribution function, but also with respect to the predictions its associated maximal entropy distribution makes on any other quantities.

Of course, although the formulae of this section are valid for any model, it requires numerical work in order to get the partition function and/or the entropy surface in a general situation.

2.2.3 The entropy and Fisher information content of Gaussian homogeneous fields

We now obtain the Shannon entropy of a family of fields when only the two-point correlation function is the relevant constraint, that we will use later in this thesis. It is easily

obtained by a straightforward generalisation of the finite dimensional multivariate case, where the means and covariance matrix of the variables are known. It is well known (Shannon, 1948) that the maximal entropy distribution is in this case the multivariate Gaussian distribution. Denoting the constraints on p_X with the matrix D and vector $\boldsymbol{\mu}$

$$\begin{aligned} D_{ij} &= \langle x_i x_j \rangle \\ \mu_i &= \langle x_i \rangle \,, \quad i,j = 1, \cdots, N \end{aligned} \tag{2.52}$$

the associated potentials are given explicitly by the relations

$$\begin{aligned} \lambda &= \frac{1}{2} C^{-1} \\ \eta &= -C^{-1}\boldsymbol{\mu}, \end{aligned} \tag{2.53}$$

where the matrix C is the covariance matrix

$$C := D - \boldsymbol{\mu}\boldsymbol{\mu}^T. \tag{2.54}$$

The Shannon entropy is given by, up to some irrelevant additive constant,

$$S(D, \boldsymbol{\mu}) = \frac{1}{2} \ln \det(D - \boldsymbol{\mu}\boldsymbol{\mu}^T). \tag{2.55}$$

The fact that about half of the constraints are redundant, due to the symmetry of the D and C matrices, is reflected by the fact that the corresponding inverse correlation matrix in equation (2.44),

$$-\frac{\partial^2 S}{\partial D_{ij} \partial D_{kl}} = -\frac{\partial \lambda_{ij}}{\partial D_{kl}} = \frac{1}{2} C^{-1}_{ik} C^{-1}_{jl}, \tag{2.56}$$

is not invertible as such if we considers all entries of the matrix D as constraints. Of course, this is not the case anymore if only the independent entries of D form the constraints.
Using the handy formalism of functional calculus, we can straightforwardly extend the above relations to systems with infinite degrees of freedom, i.e. fields, where means as well as the two-point correlation functions are constrained. A realisation of the variable X is now a field, or a family of fields $\boldsymbol{\phi} = (\phi_1, \cdots, \phi_N)$, taking values on some n-dimensional space. The expressions above in the multivariate case all stays valid, with the understanding that operations such as matrix multiplications have to be taken with respect to the discrete indices as well as the continuous ones.

With the two-point correlation function and means

$$\begin{aligned} \rho_{ij}(\mathbf{x}, \mathbf{y}) &= \langle \phi_i(\mathbf{x}) \phi_j(\mathbf{y}) \rangle \\ \bar{\phi}_i(\mathbf{x}) &= \langle \phi_i(\mathbf{x}) \rangle \end{aligned} \tag{2.57}$$

we still have, up to an unimportant constant,

$$S = \frac{1}{2} \ln \det(\rho - \boldsymbol{\phi}\boldsymbol{\phi}^T). \tag{2.58}$$

In n-dimensional Euclidean space, within a box of volume V for a family of homogeneous fields, it is simplest to work with the spectral matrices. These are defined as

$$\frac{1}{V}\left\langle \tilde{\phi}_i(\mathbf{k})\tilde{\phi}_j^*(\mathbf{k}')\right\rangle = P_{ij}(\mathbf{k})\,\delta_{\mathbf{k}\mathbf{k}'}, \tag{2.59}$$

where the Fourier transforms of the fields are defined through

$$\tilde{\phi}_i(\mathbf{k}) = \int_V d^n x\;\phi_i(x)\;e^{-i\mathbf{k}\cdot\mathbf{x}}. \tag{2.60}$$

It is well known that these matrices provide an equivalent description of the correlations, since the they form Fourier pairs with the correlation functions

$$\rho_{ij}(\mathbf{x},\mathbf{y}) = \frac{1}{V}\sum_{\mathbf{k}} P_{ij}(\mathbf{k})e^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{y})} = \rho_{ij}(\mathbf{x}-\mathbf{y}). \tag{2.61}$$

In this case, the entropy in equation (2.58) reduces, again discarding irrelevant constants, to an uncorrelated sum over the modes,

$$S = \frac{1}{2}\ln\det\left[\frac{P(0)}{V} - \bar{\phi}\bar{\phi}^T\right] + \frac{1}{2}\sum_{\mathbf{k}}\ln\det\frac{P(\mathbf{k})}{V}, \tag{2.62}$$

which is the straightforward mutlidimensional version of (Taylor and Watts, 2001, eq. 39). Comparison with equation (2.55) shows the well-known fact that the modes can be seen as Gaussian, uncorrelated and complex variables with correlation matrices proportional to $P(\mathbf{k})$. All modes have zero mean, except for the zero-mode, which, as seen from its definition, is proportional to the mean of the field itself. Accordingly, taking the appropriate derivatives, the potentials $\lambda(\mathbf{k})$ associated to $P(\mathbf{k})$ read

$$\begin{aligned}\lambda(\mathbf{k}) &= \frac{V}{2}P(\mathbf{k})^{-1}, \quad \mathbf{k}\neq 0\\ \lambda(0) &= \frac{1}{2}\left[\frac{P(0)}{V} - \boldsymbol{\phi}\boldsymbol{\phi}^T\right]^{-1}.\end{aligned} \tag{2.63}$$

and those associated to the means $\boldsymbol{\phi}$,

$$\boldsymbol{\eta} = -\left[\frac{P(0)}{V} - \boldsymbol{\phi}\boldsymbol{\phi}^T\right]^{-1}\boldsymbol{\phi} \tag{2.64}$$

Note that although the spectral matrices are, in general, complex, they are hermitian, so that the determinants are real. The amount of Fisher information in the family of fields is easily obtained with the help of equation (2.43) , with the familiar result

$$\begin{aligned}F_{\alpha\beta} = \frac{1}{2}\sum_{\mathbf{k}}\mathrm{Tr}\left[P_c^{-1}(\mathbf{k})\frac{\partial P_c(\mathbf{k})}{\partial\alpha}P_c^{-1}(\mathbf{k})\frac{\partial P_c(\mathbf{k})}{\partial\beta}\right]\\ + \frac{\partial\bar{\boldsymbol{\phi}}^T}{\partial\alpha}\left[\frac{P_c(0)}{V}\right]^{-1}\frac{\partial\bar{\boldsymbol{\phi}}}{\partial\beta},\end{aligned} \tag{2.65}$$

with $P_c(\mathbf{k})$ being the connected part of the spectral matrices,

$$P_c(\mathbf{k}) = P(\mathbf{k}) - \delta_{\mathbf{k}0} V \boldsymbol{\phi}\boldsymbol{\phi}^T. \tag{2.66}$$

These expressions are of course also valid for isotropic fields on the sphere. With a decomposition in spherical harmonics, the sum runs over the multipoles.

The Fisher matrices in common use in weak lensing or clustering can thus all be seen as special cases of this approach, namely equation (2.65), when knowledge of the statistical properties of the future data does not go beyond the two-point statistics. Indeed, in the case that the model does not predict the means, and knowing that for discrete fields the spectral matrices, equation (2.59), carry a noise term due to the finite number of galaxies, or, in the case of weak lensing, also due to the intrinsic ellipticities of galaxies, the amount of information in (2.65) is essentially identical to the standard expressions used to predict the accuracy with which parameters will be extracted from power-spectra analysis.

Of course, the maximal entropy approach, which tries to capture the relevant properties of p_X through a sophisticated guess, gives no guaranties that its predictions are actually correct. Nevertheless, as discussed in section 2.2.2, it provides a systematic approach with which to update the probability density function in case of improved knowledge of the relevant physics.

2.3 Appendix

In 2.3.1, we look at what possible 'information densities' do in fact satisfy those conditions that we would like any measure of information about the true value of a model parameter to possess. We then provide in 2.3.2 a unified derivation of the Cramér-Rao and information inequality in the multidimensional case (following a similar argumentation than in Rao (1973)) and then show its relation to maximal entropy distributions.

2.3.1 Measures of information on a parameter

Denoting with $p_X(x, \alpha)$ the probability density function of some variable, we look for functionals i

$$i[p_X(x, \alpha)] \tag{2.67}$$

such that our candidate of the measure of information on α is given by

$$I_X(\alpha) = \int dx\, i[p_X(x, \alpha)], \tag{2.68}$$

and has the following properties :

C1: We would want our functional to be a regular function of two arguments, one being the value of the probability function at α itself, and the second the value of its derivative with respect to the same α,

$$i[p_X(x, \alpha)] := i(p_X(x, \alpha), \partial_\alpha p_X(x, \alpha)). \tag{2.69}$$

While the main reason for this very strong requirement is simplicity, it appears nonetheless reasonable to us. This condition simply reflects the fact that the values of $p_X(x, \alpha)$ in regions far away from the true value of α should not provide much information, or that the information is provided by the probabilities and the linear impact of α around the true value.

C2: The choice of coordinates must not carry information, i.e. we want

$$I_X(\alpha) = I_{g(X)}(\alpha) \tag{2.70}$$

whenever g is an invertible function. While this is automatic for discrete probabilities, this is not the case for continuous distributions.

C3: The information on α from independent experiments should add

$$I_{XY}(\alpha) = I_X(\alpha) + I_Y(\alpha). \tag{2.71}$$

The last two are the key requirements for the interpretation of I as information. In combination with the first requirement, it leads to the following result: up to a multiplicative

constant, there is a unique positive definite density satisfying these conditions. it reads

$$i(p_X, \partial_\alpha p_X) = \frac{(\partial_\alpha p_X)^2}{p_X}, \tag{2.72}$$

which is easily seen to be precisely the Fisher information measure, as defined in equation (2.1).

From now on, in order to simplify the notation, we write the two arguments p_X and $\partial_\alpha p_X$ of i as p and β. If we drop the positive definite condition, the general solution is

$$i(p, \beta) = c_1 \frac{\beta^2}{p} + c_2 \frac{\beta^3}{p^2} + c_3 \beta, \tag{2.73}$$

for arbitrary constants c_{1-3}. This result follows from the claim that the most general smooth function satisfying $C1$ and $C3$ should be of the form

$$i(p, \beta) = c_1 \frac{\beta^2}{p} + c_2 \frac{\beta^3}{p^2} + c_3 \beta + c_4 \beta \ln(p) + c_5 p \ln(p), \tag{2.74}$$

for a set of arbitrary constants c_{1-5}. The two last terms do not meet $C2$ and are thus to be discarded, while the first one is the only positive definite among the three others. In order to see why this holds, we first note, as can be checked by direct calculation, that in this form $i(p, \beta)$ fulfill the requirements $C1$ and $C3$. We then use extensively $C3$ for particular instances of variables with different distributions and derivatives, sketching the proof that these criteria do not allow other functional forms.

We first notice that, since any probability density function must be normalised to unity, for any value of the model parameter, the following relations must hold for any variable X,

$$\int dx \, p_X(x, \alpha) = 1, \tag{2.75}$$

as well as

$$\int dx \, \partial_\alpha p_X(x, \alpha) = 0. \tag{2.76}$$

We now consider, beside an arbitrary distribution $p_X(x, \alpha)$, the distribution of some independent variable Y, given by $q_Y(y, \alpha)$, and denote the derivatives with respect to α associated to p and q, as f and g,

$$\begin{aligned} f(x) &= \partial_\alpha p_X(x, \alpha) \\ g(x) &= \partial_\alpha q_Y(x, \alpha), \end{aligned} \tag{2.77}$$

where the dependency on α is omitted. The joint probability density function of the independent variables X and Y is given by the product of p_X with q_Y, with associated derivative

$$\frac{\partial \, p_X q_Y}{\partial \alpha} = f q_Y + p_X g. \tag{2.78}$$

Condition $C3$ states in these terms explicitly that for any such p, q such that (2.75) holds, and f, g such that (2.76) holds, the following relation, describing the additivity of information, must be true,

$$\begin{aligned} &\int dx \int dy\, i\left(p_X(x)q_Y(y), f(x)q_Y(y) + p_X(x)g(y)\right) \\ &= \int dx\, i\left(p_X(x), f(x)\right) + \int dy\, i\left(q_Y(y), g(y)\right). \end{aligned} \tag{2.79}$$

With the help of this relation we can constrain the functional form of $i(p, \beta)$.

Let us first pick a uniform probability density for the variable Y,

$$q_Y(y) \equiv 1, \quad \text{with } y \in [0, 1]. \tag{2.80}$$

It must hold, using $C3$, that

$$\begin{aligned} &\int dx \int dy\, i\left(p_X(x), p_X(x)g(y) + f(x)\right) \\ &= \int dx\, i(p_X(x), f(x)) + \int dy\, i(1, g(y)), \end{aligned} \tag{2.81}$$

for any allowed $g(y)$. It can be shown, for instance by performing variations with respect to $g(y)$, that any solution to this integral equation must obey the relation

$$\frac{\partial^2 i}{\partial \beta^2}(p, \beta) = \frac{1}{p} j\left(\frac{\beta}{p}\right) \tag{2.82}$$

for some function j, that satisfies

$$j(\beta) = \int dx\, p_X(x) j\left(\beta + \frac{f(x)}{p_X(x)}\right). \tag{2.83}$$

This must still hold as above for any p_X and associated derivative function f, satisfying (2.75) and (2.76) respectively. Taking a variation with respect to f shows that the only solutions for $j(\beta)$ of this equation are

$$j(\beta) = c_1 \beta + c_2 \tag{2.84}$$

for some arbitrary constants c_1 and c_2. Therefore, we have constrained the full function $i(p, \beta)$ to be of the form

$$i(p, \beta) = c_1 \frac{\beta^2}{p} + c_2 \frac{\beta^3}{p^2} + \beta\, r(p) + s(p) \tag{2.85}$$

for some unknown functions $r(p)$ and $s(p)$. The first two terms are full solutions of $C2$. With very similar methods, the two other terms can be reduced to

$$r(p) = c_3 \ln(p) + c_4 \tag{2.86}$$

and

$$s(p) = c_5\, p \ln(p). \tag{2.87}$$

In this form, all terms in equation (2.85) are consistent with $C3$ and we have thus proved our claim (2.74).

2.3.2 Cramér-Rao bound and maximal entropy distributions

We denote the vector of model parameters of dimension n with

$$\boldsymbol{\alpha} = \left(\alpha_1, \cdots, \alpha_n\right) \tag{2.88}$$

and a vector of functions of dimension m the estimators

$$\hat{\mathbf{D}} = \left(\hat{D}_1, \cdots \hat{D}_m\right), \tag{2.89}$$

with expectation values $D_i(\boldsymbol{\alpha}) = \left\langle \hat{D}_i(x) \right\rangle$. In the following, we rely on Gram matrices, whose elements are defined by scalar products. Namely, for a set of vectors $\mathbf{y}_i$, the Gram matrix Y generated by this set of vectors is defined as

$$Y_{ij} = \mathbf{y}_i \cdot \mathbf{y}_j. \tag{2.90}$$

Gram matrices are positive definite and have the same rank as the set of vectors that generate them. Especially, if the vectors are linearly independent, the Gram matrix is strictly positive definite and invertible.
We adopt a vectorial notation for functions, writing scalar products between vectors as

$$f \cdot g \equiv \int dx\, p_X(x, \boldsymbol{\alpha}) f(x) g(x), \tag{2.91}$$

with $p_X(x, \boldsymbol{\alpha})$ being the probability density function of the variable X of interest. In this notation, both the Fisher information matrix and covariance matrix are seen to be Gram matrices. We have namely that the Fisher information matrix reads

$$F_{\alpha_i \alpha_j} = f_{\alpha_i} \cdot f_{\alpha_j}, \quad f_{\alpha_i}(x) = \frac{\partial \ln p_X(x, \boldsymbol{\alpha})}{\partial \alpha_i}, \tag{2.92}$$

while the covariance matrix of the estimators is

$$C_{ij} = g_i \cdot g_j, \quad g_i(x, \boldsymbol{\alpha}) = \hat{D}_i(x) - D_i(\boldsymbol{\alpha}). \tag{2.93}$$

For simplicity and since it is sufficently generic for our purpose, we will assume that both sets of vectors f and g are lineary independent, so that both matrices can be inverted. Note that we also have

$$\frac{\partial D_i}{\partial \alpha_j} = \int dx\, p_X(x, \boldsymbol{\alpha})\, \hat{D}_i(x) \frac{\partial \ln p_X(x, \boldsymbol{\alpha})}{\partial \alpha_j} = g_i \cdot f_{\alpha_j}. \tag{2.94}$$

The Gram matrix G of dimension $((m+n) \times (m+n))$ generated by the set of vectors $(g_1, \cdots, g_m, f_{\alpha_1}, \cdots, f_{\alpha_n})$ takes the form

$$G = \begin{pmatrix} C & \Delta \\ \Delta^T & F \end{pmatrix}, \quad \Delta_{i\alpha_j} = g_i \cdot f_{\alpha_j} \tag{2.95}$$

and is also positive definite due to its very definition. It is congruent to the matrix

$$YGY^T = \begin{pmatrix} C - \Delta F^{-1}\Delta^T & 0 \\ 0 & F \end{pmatrix}, \tag{2.96}$$

with

$$Y = \begin{pmatrix} 1_{m\times m} & -\Delta F^{-1} \\ 0 & 1_{n\times n} \end{pmatrix}. \tag{2.97}$$

Since two congruent matrices have the same number of positive, zero and negative eigenvalues respectively and since both F and G are positive, we can conclude that

$$C \geq \Delta F^{-1}\Delta^T, \tag{2.98}$$

which is the Cramér-Rao inequality. The lower bound on the amount of information is seen from the fact that for any matrix written in block form holds

$$\begin{pmatrix} C & \Delta \\ \Delta^T & F \end{pmatrix} \geq 0 \Leftrightarrow \begin{pmatrix} F & \Delta^T \\ \Delta & C \end{pmatrix} \geq 0 \tag{2.99}$$

and using the same congruence argument leads to the lower bound on information

$$F \geq \Delta^T C^{-1}\Delta. \tag{2.100}$$

Assume now that we have a probability density function such that this inequality is in fact an equality, i.e.

$$F = \Delta^T C^{-1}\Delta. \tag{2.101}$$

By the above argument, the Gram matrix generated by

$$(f_{\alpha_1}, \cdots, f_{\alpha_n}, g_1, \cdots, g_m) \tag{2.102}$$

is congruent to the matrix

$$\begin{pmatrix} 0_{n\times n} & 0 \\ 0 & C \end{pmatrix} \tag{2.103}$$

and has rank m. By assumption, the covariance matrix is invertible, such that the set $(g_1, \cdots, g_m)$ alone has rank m. It implies that each of the f vector can be written as linear combination of the g vectors,

$$f_{\alpha_i} = \sum_{j=1}^{m} A_j g_j, \tag{2.104}$$

or, more explicitly,

$$\frac{\partial \ln p_X(x,\boldsymbol{\alpha})}{\partial \alpha_i} = \sum_{j=1}^{m} A_j(\boldsymbol{\alpha}) \left[\hat{D}_j(x) - D_j(\boldsymbol{\alpha})\right], \tag{2.105}$$

where the key point is that the coefficients A_j are independent of x. Integrating this equation, we obtain

$$\ln p_X(x,\boldsymbol{\alpha}) = -\sum_{i=1}^{m} \lambda_i(\boldsymbol{\alpha})\hat{D}_i(x) - \ln Z(\boldsymbol{\alpha}) + \ln q_X(x) \tag{2.106}$$

for some functions λ and Z of the model parameters only, and a function q_X of x only. We obtain thus

$$p_X(x, \boldsymbol{\alpha}) = \frac{q_X(x)}{Z(\boldsymbol{\alpha})} \exp\left(- \sum_{i=1}^{m} \lambda_i(\boldsymbol{\alpha}) \hat{D}_i(x) \right). \tag{2.107}$$

This is precisely the distribution that we obtain by maximising the entropy relative to $q_X(x)$, while satisfying the constraints

$$D_i(\boldsymbol{\alpha}) = \left\langle \hat{D}_i(x) \right\rangle, \quad i = 1, \cdots, m. \tag{2.108}$$

Taking q_X as the uniform distribution makes it identical with the formula in equation (2.35).

3 Information within N-point moments

In this chapter we demonstrate how to decompose the Fisher information matrix into components unambigously associated to independent information from N-point moments of each order. The general approach to decompose the Fisher information matrix in uncorrelated components according to an orthogonal system was briefly discussed in a statistical journal in theorem 3.1 in Jarrett (1984). It seems however that this procedure was not given further attention. In this chapter, similar ideas are taken further, dealing mostly with the system of moments, where the associated orthogonal system are orthogonal polynomials.

We start in section 3.1 with one dimensional variables. We will in a first step define for a probability distribution $p(x, \boldsymbol{\theta})$, coefficients $s_n(\alpha)$ which unambiguously represent the independent information content of the moment of order n on α. These coefficients can then be used to reconstruct the Fisher information matrix order by order. The straight forward generalisation to any number of variables is performed in 3.2, or to any hierarchical system other than the N-point moments in the appendix. Properties of this expansion under the presence of noise are discussed in 3.2.2.

In section 3.3, we then present this exact decomposition for a few determinate probability density functions. We give in closed form these coefficients for several common classes of distributions, spanning a wide range of different situations. We solve these decomposition for the normal, the gamma and the beta families of distributions, as well as for an extended Poisson model.

We use this decomposition in 3.3.1 to approach the indeterminate moment problem. Namely, the Fisher information content of a moments series can be less than the information content of the probability density function it originates from, if that function cannot be uniquely recovered from the moment series. Therefore, this loss of information

is also a useful measure of the indeterminacy of the moment problem. After dealing with the lognormal distribution, whose expansion we will present in greater details in chapter 6, we treat numerically the Weibull distribution and strectched exponential. For each of those we compare the information lost to the moment hierarchy to Stoyanov's dissimilarity index of associated Stieltjes classes (Stoyanov, 2004).

3.1 One variable

Our starting point is the information inequality of chapter 2 : for any set of unbiased estimators $\hat{\mathbf{O}}$ aimed at extracting the vector of observables $\mathbf{O} = \left\langle \hat{\mathbf{O}} \right\rangle$, the following inequality between positive definite matrices holds,

$$F \geq \Delta \Sigma^{-1} \Delta^T, \tag{3.1}$$

where Σ is the covariance matrix of the estimators,

$$\Sigma_{ij} = \left\langle \hat{O}_i \hat{O}_j \right\rangle - O_i O_j, \tag{3.2}$$

and Δ is the matrix of derivatives,

$$\Delta_{\alpha i} = \frac{\partial O_i}{\partial \alpha}. \tag{3.3}$$

When the vector $\mathbf{O}$ are the moments themselves, $O_i = m_i = \left\langle x^i \right\rangle$,for $\quad i = 1, \cdots N$, it is possible to rewrite the information inequality in a more insightful form. Consider a set of polynomials $\{Q_n\}_{n=1}^{\infty}$, n the degree of the polynomial, orthogonal with respect to $p(x, \boldsymbol{\theta})$,

$$\begin{aligned} Q_n(x) &= \sum_{k=0}^{n} C_{nk} x^k \\ \langle Q_n(x) Q_m(x) \rangle &= h_n \delta_{mn}, \quad n, m = 0, \cdots, N \end{aligned} \tag{3.4}$$

where h_n is some strictly positive number, equal to unity if the polynomials are orthonormal. Since the normalisation of many common families of orthogonal polynomials is not unity, we will at the expense of extra notation keep track of the terms h_n in the following. Of course, the polynomials

$$P_n(x) := Q_n(x)/\sqrt{h_n} \tag{3.5}$$

are then orthonormal. We refer to Freud (1971) or Szegö (2003) for the theory of orthogonal polynomials. The polynomial P_n is always set by the first $2N$ moments, and is unique up to an overall sign, which we set by requiring the coefficient of x^n in $P_n(x)$ (the leading coefficient) to be positive. A simple way to build formally the orthonormal polynomials with this sign convention is for instance to apply the Gram-Schmidt orthonormalisation process to the set $(1, x, \cdots, x^n)$ with respect to the scalar product $(f, g) := \langle f(x) g(x) \rangle$.

The key point for our purposes is to realise that the inverse covariance matrix between the moments can be written [1]

$$\left[\Sigma^{-1}\right]_{ij} = \sum_{n=0}^{N} \frac{C_{in}C_{jn}}{h_n}, \quad i,j = 1, \cdots N. \tag{3.6}$$

This identity allows us to express the right hand side of the information inequality in the form of a sum of uncorrelated pieces : using (3.6), it holds

$$\sum_{i,j=1}^{N} \frac{\partial m_i}{\partial \alpha} \Sigma^{-1}_{ij} \frac{\partial m_j}{\partial \beta} = \sum_{n=1}^{N} s_n(\alpha) s_n(\beta). \tag{3.7}$$

where

$$s_n(\alpha) = \frac{1}{\sqrt{h_n}} \sum_{k=1}^{n} C_{nk} \frac{\partial m_k}{\partial \alpha}. \tag{3.8}$$

The matrix $s_N(\alpha)s_N(\beta)$ is therefore the part of the Fisher information in the Nth moment that was not contained in the moments of lower order.
These coefficients have a straightforward interpretation. They can be namely written as

$$s_n(\alpha) := \langle s(x,\alpha) P_n(x) \rangle, \quad s(x,\alpha) = \frac{\partial \ln p(x,\boldsymbol{\theta})}{\partial \alpha}, \tag{3.9}$$

which we will take as the definition of $s_n(\alpha)$. This can be seen from expanding $P_n(x)$ in terms of the matrix C, and noting that $s(x,\alpha)p(x,\boldsymbol{\theta}) = \partial_\alpha p(x,\boldsymbol{\theta})$, recovering (3.8). In other words, $s_n(\alpha)$ is nothing else than the component of the corresponding function $\partial_\alpha \ln p$ (the score function) parallel to the orthonormal polynomial of order n.

It follows immediately from equation (3.9) that the Fisher information content of the moments depends on how well the score functions can be approximated through polynomials. With increasing N, one expects the score function to be better and better reproduced by the series

$$s_{\leq N}(x,\alpha) := \sum_{n=1}^{N} s_n(\alpha) P_n(x). \tag{3.10}$$

We see that the following inequality between positive matrices holds

$$0 \leq \langle (s(x,\alpha) - s_{\leq N}(x,\alpha)) (s(x,\beta) - s_{\leq N}(x,\beta)) \rangle = F_{\alpha\beta} - \sum_{i,j=1}^{N} \frac{\partial m_i}{\partial \alpha} \left[\Sigma^{-1}\right]_{ij} \frac{\partial m_j}{\partial \beta}. \tag{3.11}$$

In particular, for any parameter α holds that the residual to the best approximation of the score function with polynomials is given by

$$\left\langle (s(x,\alpha) - s_{\leq N}(x,\alpha))^2 \right\rangle = I(\alpha) - \sum_{i,j=1}^{N} \frac{\partial m_i}{\partial \alpha} \left[\Sigma^{-1}\right]_{ij} \frac{\partial m_j}{\partial \alpha}, \tag{3.12}$$

[1] A proof can be found in a slightly more general setting later in section 3.4.2.

where $I(\alpha) = F_{\alpha\alpha}$ is the Fisher information on α. The bits of Fisher information that are absent from the set of moments m_1 to m_N are thus precisely the mean squared error of the fit of the score function through polynomials throughout the range of $p(x, \boldsymbol{\theta})$.

We define for further reference the matrices

$$[F_n]_{\alpha\beta} := s_n(\alpha)s_n(\beta) \tag{3.13}$$

as well as

$$[F_{\leq N}]_{\alpha\beta} := \sum_{n=1}^{N} s_n(\alpha)s_n(\beta) \tag{3.14}$$

These are the matrices representing the independent information content of the nth moment and of the first N moments respectively. By construction holds

$$[F_{\leq N}]_{\alpha\beta} = \sum_{i,j=1}^{N} \frac{\partial m_i}{\partial \alpha} \left[\Sigma^{-1}\right]_{ij} \frac{\partial m_j}{\partial \beta}. \tag{3.15}$$

We stress that these matrices are strictly speaking associated to a moment series rather than a density function. These matrices are namely identical for different densities having the same moment series.

We see that if the score function is itself a polynomial, of degree N, for each values of the parameters, then only the first N of these coefficients are possibly non-zero. A finite number of moments do catch all the Fisher information content of $p(x, \boldsymbol{\theta})$ in this case. In fact, the reverse statement is also true. This can be understood in the framework of orthodox statistics : in this case $p(x, \boldsymbol{\theta})$ is proportional to the exponential of a polynomial with parameter dependent coefficients, where a finite number of sufficient statistics exist Kopman (1936); Pitman and Wishart (1936). In the light of chapter 2, these distributions are precisely those that maximise Shannon entropy for fixed values of the first N moments. We thus have

$$F_{\leq N} = F \quad \text{Max. entropy distributions with constrained first } N \text{ moments.} \tag{3.16}$$

The ubiquitous example of this family being of course the Gaussian distribution, for which $N = 2$.

More generally, a sufficient criterium for the moment hierarchy to possess the same amount of Fisher information as the density function itself, is that the moment problem associated to the moment series is *determinate*, that is to say that the density can be uniquely recovered from the moment series. In this case, indeed, by a well known theorem due to M. Riesz, (Riesz, 1923), the orthogonal polynomials associated to the density form a complete set of basis functions for square integrable functions. Therefore, the square residual in equation (3.12) must go to zero for $N \to \infty$. We have in this case

$$\lim_{N\to\infty} F_{\leq N} = F \quad \text{Moment determinate density functions.} \tag{3.17}$$

Finally, for moment *indeterminate* density functions, for which different density functions exist with the same moment series, we have in general an inequality (again, an inequality between positive matrices, not matrix element to matrix element)

$$\lim_{N\to\infty} F_{\leq N} \leq F \quad \text{Moment indeterminate density functions.} \tag{3.18}$$

The amplitude of the mismatch can vary very substantially from case to case. In chapter 6, we study extensively the lognormal field, for which the inequality is always a strict inequality, but is a very strong function of the variance of the field.

3.1.1 The Christoffel-Darboux kernel as information on the density itself

To conclude this section, it is interesting to understand the interpretation of the orthogonal polynomials themselves, in this information theoretic framework. To this aim, consider that the model parameter of interests are the values of the density function themselves, allowing thus complete freedom. In the following, we do not require the density to be normalised to unity for convenience. Take for simplicity discrete values $x_i, i = 0, 1, \cdots$, with associated probability density $p(x_i)$ and the set of model parameters being precisely the density, i.e. $\alpha_i = p(x_i)$. Straight calculation leads to the following expression for the Fisher information matrix elements on the parameters $p(x_i)$, $p(x_j)$,

$$F_{ij} = \frac{\delta_{ij}}{p(x_i)}. \tag{3.19}$$

On the other hand, we have $\partial m_k / \partial p(x_i) = x_i^k$, leading to

$$s_n\left(p(x_i)\right) = P_n(x_i). \tag{3.20}$$

Therefore, the nth orthogonal polynomial $P_n(x_i)$ is the information content of the nth moment on the density function $p(x_i)$ itself. The matrices $F_{\leq N}$ become the celebrated Christoffel-Darboux kernel K_N (Simon, 2008),

$$K_N(x_i, x_j) = \sum_{n=0}^{N} P_n(x_i) P_n(x_j) = \left[F_{\leq N}\right]_{ij} \tag{3.21}$$

We have from these relations that the information escaping the first N moments is given by

$$F_{ij} - \left[F_{\leq N}\right]_{ij} = \frac{\delta_{ij}}{p(x_i)} - K_N(x_i, x_j). \tag{3.22}$$

In accordance with our discussions above, this right hand side of (3.22) is known to tend to zero as $N \to \infty$ precisely for determinate moment problem. This is can be seen from the fact that in this case, from Riesz theorem, the reproducing property holds

$$\sum_i p(x_i) \left[f(x_i) - \sum_{n=0}^{N} \left(\sum_k p(x_k) f(x_k) P_n(x_k) \right) P_n(x_i) \right]^2 \overset{N\to\infty}{\to} 0, \tag{3.23}$$

for any function $f(x_i)$ that is square summable with respect to $p(x_i)$. This implies

$$\sum_k p(x_k) f(x_k) K_N(x_k, x_i) \overset{N\to\infty}{\to} f(x_i) \tag{3.24}$$

for any such function and therefore

$$K_N(x_i, x_j) \overset{N\to\infty}{\to} \frac{\delta_{ij}}{p(x_j)}. \tag{3.25}$$

3.2 Several variables

The general theory on the statistical power of moments exposed in section 3.1 extends is a straightforward way to density functions of any number of variables and N-point moments. We first need a little bit of notation. For a d-dimensional variable X taking values $x = (x_1, \cdots x_d)$, a multiindex

$$\mathbf{n} = (n_1, \cdots, n_d), \quad n_i = 0, 1, 2, \cdots, \tag{3.26}$$

is a d-dimensional vector of non negative integers. The order of the multiindex is defined as

$$|\mathbf{n}| := \sum_{i=1}^{d} n_i. \tag{3.27}$$

For a given order $|\mathbf{n}| = n$, there are $\binom{n+d-1}{n}$ different such multiindices $\mathbf{n}$. We define further the notation $x^{\mathbf{n}}$ as

$$x^{\mathbf{n}} = x_1^{n_1} \cdots x_d^{n_d}. \tag{3.28}$$

With this notation in place, a moment of order n is given by

$$m_{\mathbf{n}} := \langle x^{\mathbf{n}} \rangle = \left\langle x_1^{n_1} \cdots x_d^{n_d} \right\rangle, \quad |\mathbf{n}| = n, \tag{3.29}$$

and the covariance matrix between the moments is

$$\langle m_{\mathbf{n}+\mathbf{m}} \rangle - m_{\mathbf{n}} m_{\mathbf{m}} =: \Sigma_{\mathbf{nm}}. \tag{3.30}$$

In this notation, the decomposition of the information in independent bits of order n proceeds by strict analogy with the one dimensional case. We refer to Dunkl and Xu (2001) for the general theory of orthogonal polynomials in several variables. A main difference being that at a fixed order N there are not one but $\binom{n+d-1}{n}$ independent orthogonal polynomials. This number is the same as the number of the above multiindices of that order N. Each multiindex defines namely an independent monomial $x^{\mathbf{n}}$ of that order. These polynomials are not defined in an unique way. The orthogonality of the

polynomials of same order is not essential for our purposes, but requiring the following condition is enough,

$$\begin{aligned} \langle P_{\mathbf{n}}(x) P_{\mathbf{m}}(x) \rangle &= 0, \quad |\mathbf{n}| \neq |\mathbf{m}| \\ \langle P_{\mathbf{n}}(x) P_{\mathbf{m}}(x) \rangle &= [H_n]_{\mathbf{nm}}, \quad |\mathbf{n}| = |\mathbf{m}| = n \end{aligned} \tag{3.31}$$

for some positive matrices H_n, which replace the normalisation h_n in 3.1. The component of the score function $\partial_\alpha \ln p$ parallel to the polynomial $P_{\mathbf{n}}$ is

$$s_{\mathbf{n}}(\alpha) := \left\langle \frac{\partial \ln p(x, \boldsymbol{\theta})}{\partial \alpha} P_{\mathbf{n}}(x) \right\rangle, \tag{3.32}$$

and the expansion of the score function in terms of these polynomials reads

$$s_{\leq N}(x, \alpha) := \sum_{n=0}^{N} \sum_{|\mathbf{n}|,|\mathbf{m}|=n} s_{\mathbf{n}}(\alpha) \left[H_n^{-1}\right]_{\mathbf{nm}} P_{\mathbf{m}}(x). \tag{3.33}$$

It will converge to the score function for $N \to \infty$ if the set of polynomials is complete, whereas it may not if not. We note that there is some freedom in the definition (3.31). This freedom is that of the choice of a basis in the vector space of polynomials of order n orthogonal to all polynomials of lower order. For this reason, $s_{\mathbf{n}}(\alpha)$ depends on the particular basis. However, the expansion (3.33) does not, and so will not the information matrices at fixed order.

Writing now the orthogonal polynomials in terms of a triangular transition matrix

$$P_{\mathbf{n}}(x) = \sum_{|\mathbf{m}| \leq |\mathbf{n}|} C_{\mathbf{nm}} x^{\mathbf{m}} \tag{3.34}$$

we see that the information matrix of order n (3.13) becomes

$$\begin{aligned} [F_n]_{\alpha\beta} &= \sum_{|\mathbf{n}|,|\mathbf{m}|=n} s_{\mathbf{n}}(\alpha) \left[H_n^{-1}\right]_{\mathbf{nm}} s_{\mathbf{m}}(\beta) \\ &= \sum_{|\mathbf{i}|,|\mathbf{j}|\leq n} \left[C^T H_n^{-1} C\right]_{\mathbf{ij}} \frac{\partial m_{\mathbf{i}}}{\partial \alpha} \frac{\partial m_{\mathbf{j}}}{\partial \alpha}, \end{aligned} \tag{3.35}$$

The strict analog of equation (3.14) holds for each N,

$$[F_{\leq N}]_{\alpha\beta} = \sum_{n=1}^{N} [F_n]_{\alpha\beta} = \sum_{|\mathbf{i}|,|\mathbf{j}|=1}^{N} \frac{\partial m_{\mathbf{i}}}{\partial \alpha} \left[\Sigma^{-1}\right]_{\mathbf{ij}} \frac{\partial m_{\mathbf{j}}}{\partial \beta}, \tag{3.36}$$

recovering the right hand side of the information inequality for all moments of order up to N. Just as before, the missing piece between F and $F_{\leq N}$ is the least squared residual the approximation of the score function through polynomials of order up to N.
The matrices F_n and $F_{\leq N}$ are easily seen to be invariant under mappings $y = Ax + b$, where A is an invertible square matrix of size $d \times d$ and b a d dimensional vector, provided both are parameter independent.

As a simple illustration, for the multivariate Gaussian distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix C we have

$$[F_1]_{\alpha\beta} = \frac{\partial \boldsymbol{\mu}}{\partial \alpha} C^{-1} \frac{\partial \boldsymbol{\mu}}{\partial \beta} \tag{3.37}$$

and

$$[F_2]_{\alpha\beta} = \frac{1}{2} \mathrm{Tr} \left[C^{-1} \frac{\partial C}{\partial \alpha} C^{-1} \frac{\partial C}{\partial \beta} \right], \tag{3.38}$$

summing up to the total information, $F_n = 0, n > 2$.

3.2.1 Independent variables and uncorrelated fiducial

In general, it is a rather difficult problem to obtain explicit expressions for the orthogonal polynomials or the matrices F_N and $F_{\leq N}$, especially in the case of several variables. Using exact though formal expressions as starting point requires the evaluation of determinants of moment matrices, and cases are rare when this is tractable. For independent variables $(x_1, \cdots, x_d)$, however, a canonical choice of $P_{\mathbf{n}}$ can be written down in terms of the ones of associated to one dimensional problem. More specifically, if the variables are independent, then

$$p(x, \boldsymbol{\theta}) = \prod_{i=1}^{d} p_i(x_i, \boldsymbol{\theta}), \tag{3.39}$$

where p_i denotes the one dimensional probability density function of the ith variable. Define then the polynomial in d variables of order $|\mathbf{n}|$ through

$$P_{\mathbf{n}}(x) := \prod_{i=1}^{d} P_{n_i}(x_i), \tag{3.40}$$

where $P_{n_i}(x_i)$ is the orthonormal polynomial in one variable of order i with respect to p_i. It is not difficult to see that for any multiindices $\mathbf{n}$ and $\mathbf{m}$ the average over x factorizes in averages with respect to each variable x_i. We have namely

$$\langle P_{\mathbf{n}}(x) P_{\mathbf{m}}(x) \rangle = \prod_{i=1}^{d} \langle P_{n_i}(x_i) P_{m_i}(x_i) \rangle = \prod_{i=1}^{d} \delta_{n_i m_i} = \delta_{\mathbf{nm}}. \tag{3.41}$$

The so defined polynomials are therefore orthogonal with matrices H_n being unit matrices. From equation (3.39) follows that the function $\partial_\alpha \ln p$ is the sum of the functions $\partial_\alpha \ln p_i$ of the individual variables. Therefore, using the above polynomials defined in (3.40), it is not difficult to see that all the coefficients $s_{\mathbf{n}}(\alpha)$ that couples different variables vanishes, i.e. $s_{\mathbf{n}}(\alpha) = 0$ if $\mathbf{n}$ has two or more non zero indices. The information content of order N becomes then simply the sum of the information of order N within each variable, as expected. This is an manifestation of the additivity of Fisher information for independent variables, which is thus seen to hold order by order.

Note that polynomials also take this product form (3.40) if as above the probability density function factorizes at the fiducial values of the model parameters, but with the difference that there is no splitting of the derivative functions $\partial_\alpha \ln p$ as a uncorrelated sum. In this case, the fiducial model is uncorrelated, but model parameters create correlations away from their fiducial values. N-point moments at non zero lag may carry genuine information in this case. In mathematical terms, all $s_{\mathbf{n}}(\alpha)$ can possibly be non zero, depending on how the score functions couple the different variables.

3.2.2 On the impact of noise

Very often the observed random variable is the sum of two random variables,

$$z = x + y, \tag{3.42}$$

where x is a d-dimensional variable carrying the information on the parameters $\boldsymbol{\theta}$, and y is some additive noise. We set

$$p_X(x, \boldsymbol{\theta}) \tag{3.43}$$

to be the probability density function for x, and

$$p_Y(y) \tag{3.44}$$

being that for y, which is parameter independent. From the rules of probability theory, we have

$$p_Z(z, \boldsymbol{\theta}) = \int d^d x \, p_X(x, \boldsymbol{\theta}) \, p_Y(z - x). \tag{3.45}$$

In this section, we prove the following, maybe rather intuitive fact, valid for any functional form of p_X and of p_Y, any model parameters and any dimensionality d:
For any N, the following relation between positive matrices holds

$$F^{X+Y}_{\leq N} \leq F^{X}_{\leq N}. \tag{3.46}$$

On the other hand, this relation does not hold for the matrices F_n. For instance, a non Gaussian noise Y on a Gaussian signal X will create in general third and higher order terms in the score functions. However, their amplitude is constrained by (3.46).

We prove this inequality here only for one dimensional variables, $d = 1$. The general proof consists in replacing indices such as n with multindices $\mathbf{n}$ and so on.

Our proof is based on the following basic fact concerning blockwise positive matrices (for a proof of this particular fact, see (Bhatia, 2007, e.g.) : Whenever A and D are strictly positive matrices, then

$$\begin{pmatrix} A & C \\ C^T & D \end{pmatrix} \geq 0 \Leftrightarrow A \geq C D^{-1} C^T. \tag{3.47}$$

We proceed as follows : in general we can write the moments of Z as a linear combination of those of X

$$m_n^Z = \langle (x+y)^n \rangle = \sum_{k=0}^{n} \binom{n}{k} m_n^X m_{n-k}^Y =: \sum_{k=0}^{n} B_{nk} m_k^X. \tag{3.48}$$

Note that B is a triangular matrix with non vanishing diagonal elements, and is therefore invertible. Writing $\mathbf{m} = (m_0, \cdots m_n)^T$, it holds therefore

$$\left[F_{\leq N}^Z\right]_{\alpha\beta} = \left(\frac{\partial \mathbf{m}^Z}{\partial \alpha}\right)^T M_Z^{-1} \left(\frac{\partial \mathbf{m}^Z}{\partial \beta}\right) = \left(\frac{\partial \mathbf{m}^X}{\partial \alpha}\right)^T B^T M_Z^{-1} B \left(\frac{\partial \mathbf{m}^X}{\partial \beta}\right), \tag{3.49}$$

where M_X, respectively M_Z is the moment matrix $M_{ij} = m_{i+j}$ of X, respectively Z. The claim (3.46) is proven provided

$$\left(\frac{\partial \mathbf{m}^X}{\partial \alpha}\right)^T B^T M_Z^{-1} B \left(\frac{\partial \mathbf{m}^X}{\partial \beta}\right) \leq \left(\frac{\partial \mathbf{m}^X}{\partial \alpha}\right)^T M_X^{-1} \left(\frac{\partial \mathbf{m}^X}{\partial \beta}\right), \tag{3.50}$$

which is equivalent to

$$B^T M_Z^{-1} B \leq M_X^{-1}. \tag{3.51}$$

We now prove this last relation.

Consider the following vector

$$\mathbf{v} := (z^0, \cdots, z^N, x^0, \cdots, x^N) \tag{3.52}$$

and the matrix G defined with the help of this vector,

$$G_{ij} := \langle v_i v_j \rangle := \int dx \int dy\, p_X(x) p_Y(y) v_i v_j. \tag{3.53}$$

Such a matrix of scalar products is called a Gram matrix. G has by definition the block form

$$G = \begin{pmatrix} M_Z & B M_X \\ M_X B^T & M_X \end{pmatrix}. \tag{3.54}$$

It is well known that any Gram matrix is positive : from the definition of G, we have that $u^T G u = \left\langle (u \cdot v)^2 \right\rangle \geq 0$ for any vector u. Using the above fact (3.47) we have

$$M_Z \geq B M_X B^T. \tag{3.55}$$

Equivalently

$$M_X \leq B^{-1} M_Z B^{-T}. \tag{3.56}$$

Taking the inverse gives

$$M_X^{-1} \geq B^T M_Z^{-1} B, \tag{3.57}$$

which concludes the proof.

3.2.3 Poissonian discreteness effects

The same relations hold for another source of noise relevant in cosmological surveys, i.e. discreteness effects due to a finite number of tracers of the underlying fields. A common parametrisation is the Poisson model, where the observed number of tracers in a cell is given by a Poisson variable with intensity the value of the underlying continuous field x at that point. Explicitly, with $N = (N_1, \cdots, N_d)$ the number of tracers in d cells, we set their joint probability to be

$$p_N(N_1, \cdots, N_d) = \int d^d x \, p_X(x, \boldsymbol{\theta}) \prod_{i=1}^{d} e^{-x_i} \frac{x_i^{N_i}}{N_i!}. \tag{3.58}$$

This model has the peculiar property of transforming the moments of X to factorial moments : the falling factorial in d variables[2] becomes

$$(N)_{\mathbf{m}} = \prod_{i=1}^{d} (N_i)_{m_i} \tag{3.59}$$

and we have indeed the known curious relation

$$\left\langle N_{(\mathbf{m})} \right\rangle = \langle x^{\mathbf{m}} \rangle = m^X_{\mathbf{m}}. \tag{3.60}$$

Using this relation, it is easy to prove with the same methods than above that

$$F^N_{\leq M} \leq F^X_{\leq M} \tag{3.61}$$

for any M. This does not appear to be necessarily the case for more generic functional form of $p_N(N|x)$, since there are no obvious relations between the first M moments of N and those of X.

3.3 Some exactly solvable models for moment determinate densities

We derive in this section the exact analytical expressions for the information coefficients at all orders of well known families of moment determinate probability density functions. We will deal with the normal distribution, the beta and gamma families as well as an extended Poisson model. For all these instances, the matrices $F_{\leq N}$ converges to F as $N \to \infty$.

Normal distribution The normal distribution provides us with a very simple illustration of the approach. Its probability density function is

$$p(x, \mu, \sigma^2) := \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2} \frac{(x-\mu)^2}{\sigma^2} \right). \tag{3.62}$$

[2] see (3.95) for the falling factorial in one dimension

Its Fisher information matrix is well known,

$$F = \begin{pmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{1}{2\sigma^4} \end{pmatrix}. \tag{3.63}$$

The information coefficients take the form

$$s_1(\mu) = \frac{1}{\sigma}, \quad s_n(\mu) = 0, \quad n \neq 1, \tag{3.64}$$

and

$$s_2(\sigma^2) = \frac{1}{\sqrt{2}\sigma^2}, \quad s_n(\sigma^2) = 0, \quad n \neq 2. \tag{3.65}$$

We recover already the full matrix with the first two moments, $F_{\leq 2} = F$. The matrix is diagonal because there is no order n for which both coefficients $s_n(\mu)$ and $s_n(\sigma^2)$ do not vanish.

Derivation : It is easily seen that the score function $\partial_\mu \ln p$ is a polynomial of first order in x. This implies immediately that only $s_1(\mu)$ is non-zero. For any probability density function and parameter α, it holds from (3.8) that

$$s_1(\alpha) = \frac{\partial m_1}{\partial \alpha} \frac{1}{\sigma}, \tag{3.66}$$

which proves (3.64). Very similarly, the score function associated to σ^2 is a polynomial of second order, such that only the first two coefficients can possibly be non-zero. However, $s_1(\sigma^2)$ vanishes since σ^2 does not impact the mean. Finally, $s_2(\sigma^2)$ can be gained by noting that the polynomials orthogonal to the normal distribution are the Hermite polynomials (Szegö, 2003),

$$P_n(x) = H_n\left(\frac{x-\mu}{\sigma}\right), \tag{3.67}$$

with normalisation $h_n = 1/n!$. From $H_2(x) = x^2 - 1$ and equation (3.8) follows $C_{22} = 1/\sigma^2$ and $s_2(\sigma^2) = \frac{1}{\sqrt{2}\sigma^2}$.

Beta distribution The beta distribution is defined as

$$p(x, \alpha, \beta) := \frac{x^{\alpha-1}\,(1-x)^{\beta-1}}{B(\alpha, \beta)}, \quad 0 < x < 1, \;\; \alpha, \beta > 0 \tag{3.68}$$

where $B(\alpha, \beta)$ is the beta integral, which will also enter the following section on the gamma distribution. It has the well known representation in terms of the gamma function Γ, $B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$. It plays a fundamental role in order statistics.

The full Fisher information matrix can be conveniently expressed in terms of the second derivative of the logarithm of the gamma function, called the trigamma function ψ_1 (Abramowitz and Stegun, 1970, p. 258-260),

$$F = \begin{pmatrix} \psi_1(\alpha) & 0 \\ 0 & \psi_1(\beta) \end{pmatrix} - \psi_1(\alpha + \beta) \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \tag{3.69}$$

Our derivation of the information coefficients associated to α and β, based on the explicit expressions of the orthogonal polynomials, in this case the Jacobi polynomials, is rather lengthy. We defer it to the appendix.

The result is

$$s_n(\alpha) = -\frac{(-1)^n}{n}\left(\frac{(\beta)^{(n)}n!}{(\alpha)^{(n)}(\alpha+\beta)^{(n)}}\right)^{1/2}\left(\frac{2n+\alpha+\beta-1}{n+\alpha+\beta-1}\right)^{1/2}. \tag{3.70}$$

and

$$s_n(\beta) = -\frac{1}{n}\left(\frac{(\alpha)^{(n)}n!}{(\beta)^{(n)}(\alpha+\beta)^{(n)}}\right)^{1/2}\left(\frac{2n+\alpha+\beta-1}{n+\alpha+\beta-1}\right)^{1/2}, \tag{3.71}$$

where $(x)^{(n)} := x(x+1)\cdots(x+n-1)$ is the rising factorial. For the symmetric case $\alpha = \beta = 1/2$, these expressions simplify considerably. We have namely

$$s_n(\alpha) = (-1)^{n-1}\frac{\sqrt{2}}{n}, \quad s_n(\beta) = -\frac{\sqrt{2}}{n}. \tag{3.72}$$

Since $\psi_1(1/2) = \pi^2/2$ and $\psi_1(1) = \pi^2/6$, a short calculation shows that we recover indeed, summing these coefficients, the full matrix given in (3.69)

$$F|_{\alpha=\beta=1/2} = \frac{\pi^2}{6}\begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}. \tag{3.73}$$

Gamma distribution The gamma distribution is a two parameter family,

$$p(x, \ln\theta, k) := \frac{e^{-x/\theta}}{\theta^k\Gamma(k)}x^{k-1} \quad k, \theta > 0, \quad x > 0. \tag{3.74}$$

where $\Gamma(k)$ is the gamma function. Special cases include the exponential distribution ($k = 1$), or the chi squared distribution with n degrees of freedom ($k = n/2, \theta = 2$). The calculation of the Fisher information matrix associated to $\ln\theta$ and k is not difficult. Again, the trigamma function shows up,

$$F = \begin{pmatrix} k & 1 \\ & \\ 1 & \psi_1(k) \end{pmatrix}. \tag{3.75}$$

The information coefficients are evaluated below, with the result

$$s_1(\ln\theta) = \sqrt{k}, \quad s_n(\ln\theta) = 0, \quad n \neq 1, \tag{3.76}$$

and

$$s_n(k) = (-1)^{n-1}\sqrt{\frac{B(k,n)}{n}}. \tag{3.77}$$

As a consistency check, we see that since $s_1(k) = \frac{1}{\sqrt{k}}$, we recover trivially the $\ln\theta\ln\theta$ and $\ln\theta k$ elements of the matrix (3.75). Its kk element implies that the following identity must hold,

$$\sum_{n=1}^{\infty} \frac{B(k,n)}{n} = \psi_1(k), \quad k > 0, \tag{3.78}$$

which reduces for $k = 1$ to Euler's famous formula $\sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{6}$.

Derivation : The polynomials associated to that distribution are the generalized Laguerre polynomials $L_n^{(k-1)}$(Abramowitz and Stegun, 1970, p.775) . More precisely, we have

$$\int_0^{\infty} t^{k-1} e^{-t} L_n^{(k-1)}(t) L_m^{(k-1)}(t) = \frac{\Gamma(n+k)}{n!}\delta_{mn}. \tag{3.79}$$

For this reason, the polynomials orthogonal to the gamma distribution are

$$P_n(x) := L_n^{(k-1)}\left(\frac{x}{\theta}\right), \tag{3.80}$$

with normalisation

$$h_n = \frac{\Gamma(n+k)}{\Gamma(k)}\frac{1}{n!} = \frac{1}{n}\frac{1}{B(k,n)}. \tag{3.81}$$

They have the explicit matrix elements C_{ni}[3],

$$C_{ni} = (-1)^{n-i}\binom{n+k-1}{n-i}\frac{\theta^{-i}}{i!}. \tag{3.82}$$

The moments of the gamma function are given by

$$m_n = \theta^n \frac{\Gamma(n+k)}{\Gamma(k)}. \tag{3.83}$$

The Fisher information on θ is the simplest. It holds namely that the score function associated to θ is a polynomial first order in x, and therefore that only $s_1(\ln\theta)$ is non-zero. It follows

$$s_1(\ln\theta) = \sqrt{k}, \quad s_n(\ln\theta) = 0, \quad n \neq 1. \tag{3.84}$$

The calculation of $s_n(k)$ requires a little bit more work, but we can make use of previous results we derived when dealing with the beta family. The derivatives of the moments with respect to k are given by

$$\frac{\partial \ln m_n}{\partial k} = \psi_0(k+n) - \psi_0(k). \tag{3.85}$$

Using the representation (3.8), together with

$$\binom{n+k-1}{n-i}\frac{1}{i!}\frac{1}{\Gamma(k+i)} = \frac{\Gamma(n+k)}{\Gamma(n+1)}\binom{n}{i}. \tag{3.86}$$

[3]We added the factor $(-1)^n$ to the conventions of Abramowitz and Stegun (1970), in accordance with our own conventions of having a positive leading coefficient C_{nn}.

one obtains the following expression

$$s_n(k) = \frac{1}{\sqrt{n\, B(k,n)}} \sum_{i=0}^{n} (-1)^{n-i} \binom{n}{i} \left(\psi_0(k+i) - \psi_0(k)\right). \tag{3.87}$$

The sum is precisely the function $f_q(t)$ defined in (3.142), evaluated a $t = 1$ and $k = q$, and that we proved in (3.150) to be a representation of the beta function,

$$f_k(1) = (-1)^{n-1} B(k,n). \tag{3.88}$$

We conclude therefore

$$s_n(k) = (-1)^{n-1} \sqrt{\frac{B(k,n)}{n}}. \tag{3.89}$$

Poisson model Consider the probability density function

$$p(x, \mu, \ln A) = \sum_{n=0}^{\infty} e^{-\mu} \frac{\mu^n}{n!} \delta^D(x - An), \quad \mu, A > 0 \tag{3.90}$$

with δ^D the Dirac delta function. This is the usual Poisson law, together with some amplitude A that is left free. If A is set to unity, we recover the Poisson law. This is a peculiar situation in terms of Fisher information. Unlike the previous situations, the parameter A impacts indeed the range of x. The Fisher information on A is not well defined, formally infinite.

Even though the Fisher information matrix is not well defined, the information coefficients as given in equations (3.8) or (3.9) are still meaningful, and so are their interpretation (3.7) in terms of the expected curvature of a χ^2 fit. The result is

$$s_1(\ln \mu) = \sqrt{\mu}, \quad s_n(\ln \mu) = 0, \quad n \neq 1. \tag{3.91}$$

and

$$s_1(\ln A) = \sqrt{\mu}, \quad s_n(\ln A) = (-1)^n \frac{\sqrt{n!}}{n(n-1)} \mu^{1-n/2}, \quad n \geq 2. \tag{3.92}$$

Note that it is compatible with the fact that the Fisher information on A is infinite. For any μ the sum

$$\sum_{n=1}^{\infty} s_n^2(\ln A) \tag{3.93}$$

is divergent. The turnover of the coefficients occurs around $n \approx \mu$ where they start to increase. Atypically, the higher order the moment the most interesting it becomes for inference on A in this model.

Derivation : Consider the Charlier polynomials (Abramowitz and Stegun, 1970, pp.788), defined by

$$c_n(x) := \sum_{k=0}^{n} (-1)^{n-k} \binom{n}{k} \mu^{-k} (x)_k, \tag{3.94}$$

where

$$x(x-1)\cdots(x-m+1) =: (x)_m \tag{3.95}$$

is the falling factorial. These are the polynomials orthogonal with respect to the Poisson distribution with intensity μ. Clearly, the polynomials

$$P_n(x) := c_n\left(\frac{x}{A}\right) \tag{3.96}$$

are then the polynomials orthogonal to $p(x, \mu, \ln A)$. The normalization is

$$h_n = \frac{n!}{\mu^n}. \tag{3.97}$$

On the other hand, the moments are given by $m_n = A^n m_n^P$, where m_n^P is the nth moment of the Poisson distribution. For this reason, it holds that $\frac{\partial m_n}{\partial \ln A} = n\, m_n$. This equation, together with (3.96) implies that the information coefficients on $\ln A$ are independent of the actual value of A, since $C_{nk} \propto A^{-k}$ and $m_k \propto A^k$. So are the coefficients $s_n(\ln \mu)$. We can therefore from now on safely chose $A = 1$ as the fiducial value, and work with the usual Poisson distribution and associated Charlier polynomials.

The coefficients associated to $\ln \mu$ are the simplest to obtain. It is not difficult to see that the score function associated to $\ln \mu$ of the Poisson distribution is a polynomial of first order. We obtain $s_1(\ln \mu) = \sqrt{\mu}$ and $s_n(\ln \mu) = 0, \quad n \neq 1$. Turning to $\ln A$, we need to evaluate

$$s_n(\ln A) = \frac{\mu^{n/2}}{\sqrt{n!}} \sum_{k=0}^{n} C_{nk}\, k\, m_k. \tag{3.98}$$

The derivation we propose uses extensively the technique of generating functions (Wilf, 1994), and is a variation on the following theme : consider the known expression for the factorial moments of the Poisson distribution,

$$\langle n(n-1)\cdots(n-m+1)\rangle = \mu^m. \tag{3.99}$$

Even though it may appear mysterious at first sight, a simple way to prove this identity is to notice

$$\langle n(n-1)\cdots(n-m+1)\rangle = \frac{d^m}{dt^m} G(t)\bigg|_{t=1}, \quad G(t) := \sum_{n=0}^{\infty} p(n) t^n, \tag{3.100}$$

where $G(t)$ is called the probability generating function, and to try and evaluate the right hand side. This is indeed very convenient, since the probability generating function of the poisson law takes a very simple form,

$$G(t) = e^{-\mu} \sum_{n=0}^{\infty} \frac{\mu^n}{n!} t^n = e^{\mu(t-1)}. \tag{3.101}$$

The derivatives are trivial to evaluate, proving (3.99). We will need the following result, that comes out directly from an essentially identical argument,

$$\langle (an)_m \rangle = \frac{d^m}{dt^m} e^{\mu(t^a-1)}\bigg|_{t=1}, \tag{3.102}$$

where a is some number.

We now turn to the evaluation of

$$\sum_{k=0}^{n} C_{nk} k\, m_k. \tag{3.103}$$

This sum can be written as the first derivative with respect to a evaluated at $a = 1$ of the following expression

$$\langle c_n(ax) \rangle = \sum_{k=0}^{\infty} \frac{e^{-\mu}\mu^k}{k!} c_n(ak), \tag{3.104}$$

as can be seen directly from an expansion of the Charlier polynomials c_n in terms of the coefficients C_{nk}. Using equation (3.94) and our result (3.102), we obtain

$$\sum_{k=0}^{n} C_{nk} k\, m_k = e^{-\mu} \sum_{k=0}^{n} (-1)^{n-k} \mu^{-k} \binom{n}{k} \frac{d}{da} \frac{d^k}{dt^k} \exp(\mu t^a)\bigg|_{t,a=1}. \tag{3.105}$$

The derivative with respect to a is easily performed :

$$\frac{d}{da} \exp(\mu t^a)\bigg|_{a=1} = \mu\, e^{\mu t}\, t \ln t. \tag{3.106}$$

Also, using Leibniz rule for derivatives of products, we can write

$$\frac{d^k}{dt^k} \exp(\mu t)\ t \ln t\bigg|_{t=1} = e^{\mu} \sum_{i=0}^{k} \binom{k}{i} \mu^{k-i}\ \frac{d^i}{dt^i}\, t \ln t\bigg|_{t=1}. \tag{3.107}$$

The factor of μ^k cancels in equation (3.105), so that

$$\sum_{k=0}^{n} C_{nk} k\, m_k = \mu \sum_{k=0}^{n} \sum_{i=0}^{k} \binom{n}{k} \binom{k}{i} (-1)^{n-k} \mu^{-i}\ \frac{d^i}{dt^i}\, t \ln t\bigg|_{t=1}. \tag{3.108}$$

From the properties of the binomial coefficients follows that the sum over k is just $(-1)^n \delta_{ni}$. Therefore,

$$s_n(\ln A) = \mu \frac{\mu^{-n/2}}{\sqrt{n!}}\ \frac{d^n}{dt^n}\, t \ln t\bigg|_{t=1}. \tag{3.109}$$

The derivatives are easily computed :

$$\frac{d^n}{dt^n}\, t \ln t\bigg|_{t=1} = \begin{cases} 0 & n = 0 \\ 1 & n = 1 \\ (-1)^n (n-2)! & n \geq 2 \end{cases} \tag{3.110}$$

with the final result given in equation (3.92).

3.3.1 Moment indeterminate densities

We now turn to indeterminate distributions, i.e. those that cannot be uniquely recovered from their moments. In this case, the Fisher information content of the moments will generally not converge to the total amount. In this respect, the limit

$$\lim_{N\to\infty} \frac{\sum_{n=1}^{N} s_n^2(\alpha)}{I(\alpha)}, \tag{3.111}$$

a number in $(0, 1]$, may be thought of as an indirect measure of the amount of indeterminacy of the moment problem. A value of unity means that the moments do carry the same information on the parameter α as the probability density function, while for values close to zero a large amount of information on α is not present in the hierarchy. The quantity $\sqrt{\epsilon(\alpha)}$ has via the information inequality the interpretation of the ratio of the expected constraints on α by extraction of the full set of moments to the best constraints achievable with the help of unbiased estimators of α. When more than one parameter are of interest, this definition is not really satisfactory anymore, since it is tied to α exclusively. One way to get around this is by defining a ratio of determinants,

$$\epsilon := \lim_{N\to\infty} \epsilon_N = \lim_{N\to\infty} \det F_{\leq N}/\det F, \tag{3.112}$$

where $F_{\leq N}$ was defined in equation (3.14). Clearly, ϵ_N is the cumulative efficiency to catch the information. Since the matrices in the numerator and denominator do transform in the same way under any smooth reparametrisation of the parameters, or coordinates, of the family, this ratio is left unchanged under such transformations. Thus, ϵ and ϵ_N are well defined quantities associated to that family of distributions, independently of the chosen coordinates.

Naively, one might expect this lost information to reflect in some way the freedom there is in the choice of a distribution with the very same series of moments. We will therefore compare our results to a measure of this freedom, D_S, that was proposed recently (Stoyanov, 2004) as the index of dissimilarity of a Stieltjes class. A Stieltjes class is a family of probability density functions of the form

$$p_\delta(x) = p(x)\left(1 + \delta h(x)\right), \quad |h(x)| \leq 1, \ |\delta| \leq 1, \tag{3.113}$$

all having the same moment series as the central probability density function p. The index D_S, defined as

$$\int dx\, p(x)|h(x)|, \tag{3.114}$$

in $[0, 1]$ is the maximal distance between two members of the class, being zero for determinate moment problems, where h must identically vanish.

We will consider three types of probability density functions, the lognormal, which is indeterminate for any values of its parameter space, as well as the Weibull distribution and the stretched exponential, which are indeterminate solely for parts of their parameter range. For these last two distributions we could not find in all cases exact analytical

expressions for the coefficients s_n. We generated them therefore numerically, using a fine discretization procedure to obtain the orthogonal polynomials, following the exposition in Gautschi (2004).

Lognormal distribution The lognormal distribution has the following functional form,

$$p(x, \mu_Y, \sigma_Y) = \frac{1}{\sqrt{2\pi}\sigma_Y} \frac{1}{x} \exp\left(-\frac{1}{2} \frac{(\ln x - \mu_Y)^2}{\sigma_Y^2} \right), \quad x > 0 \tag{3.115}$$

The parameters μ_Y and σ_Y^2 are the mean and the variance of $y = \ln x$ which is normally distributed. They are related to the mean μ and variance σ^2 of x by the relations

$$\begin{aligned} \mu_Y &= \ln \mu - \frac{1}{2} \ln\left(1 + \frac{\sigma^2}{\mu^2} \right) \\ \sigma_Y^2 &= \ln\left(1 + \frac{\sigma^2}{\mu^2} \right). \end{aligned} \tag{3.116}$$

The only relevant parameter for our purposes is the reduced variance σ^2/μ^2, in terms of which our results can be expressed. The information coefficients associated to the parameters $\ln \mu$ and σ^2/μ^2 we will derive in details in chapter 6, (see also Carron (2011)), with the help of q-series. These are given by

$$\begin{aligned} s_n(\ln \mu) &= (-1)^{n-1} \left(\frac{q^n}{1-q^n} \right)^{1/2} \sqrt{(q:q)_{n-1}} \\ s_n(\sigma^2/\mu^2) &= q(-1)^n \left(\frac{q^n}{1-q^n} \right)^{1/2} \sqrt{(q:q)_{n-1}} \left(\sum_{k=1}^{n-1} \frac{q^k}{1-q^k} \right). \end{aligned} \tag{3.117}$$

In these equations,

$$q := \frac{1}{1 + \sigma^2/\mu^2} \tag{3.118}$$

and

$$(q:q)_n := \prod_{k=1}^{n} \left(1 - q^k \right) \tag{3.119}$$

is the q-Pochammer symbol Andrews et al. (1999). According to the chain rule of derivation, we obtain $s_n(\mu_Y) = s_n(\ln \mu)$ and $s_n(\sigma_Y^2) = \frac{1}{2}(s_n(\ln \mu) + s_n(\sigma^2/\mu^2))$. On the other hand, since $\ln x$ is normally distributed, the total Fisher information matrix F associated to μ_Y and σ_Y^2 is diagonal with $1/\sigma_Y^2$ and $1/2\sigma_Y^4$ as diagonal elements. It follows that in these coordinates $\det F = 1/(2\sigma_Y^6)$.

The Stieltjes class we consider is set by

$$h(x) = \sin\left[\frac{2\pi}{\sigma_Y^2} (\ln x - \mu_Y) \right]. \tag{3.120}$$

Others are given in Heyde (1963) but are equivalent to a uninteresting rescaling of σ_Y. After an obvious variable substitution, the dissimilarity index becomes

$$D_S = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\frac{x^2}{2}\right) \left|\sin\left(\frac{2\pi}{\sigma_Y} x\right)\right| \, dx. \tag{3.121}$$

In figure 3.1, we show the cumulative efficiency ϵ_N, for $N = 2, 5, 10, 15$ and 1000 for

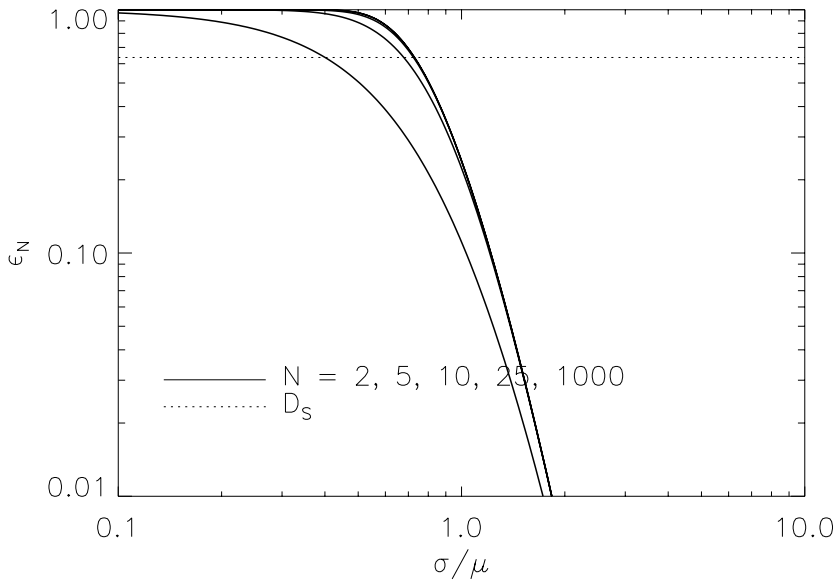

Fig. 3.1 — Solid lines : the cumulative efficiencies ϵ_N, for $N = 2, 5, 10, 25$ and 1000, from bottom to top, as function of the reduced variance of the lognormal distribution. Dashed : The dissimilarity index for the Stieltjes class given in equation (3.120).

the lognormal family, evaluated with the exact expressions given in (3.117), solid lines from bottom to top, together with the dissimilarity index D_S, shown as the dashed line, evaluated by numerical integration of equation (3.121). It is not difficult to see from the above expressions that the coefficients s_n decay exponentially for large N, such that the convergence is rather quick over the full range. We observe a very sharp transition in ϵ as soon as the reduced variance approaches unity, from a regime where the entire information content is within the second moment, as for the normal distribution (and thus where the indeterminacy of the moment problem is irrelevant for parameter estimation), to the opposite regime where all moments completely fail to capture the information. This cutoff is discussed at length in chapter 6. On the other hand. the index D_S is seen to be remarkably constant over the range shown, equal to its limiting values for $\sigma_Y \to 0$, which can be evaluated from (3.121) to be $2/\pi$.

Weibull distribution We consider now the Weibull distribution with shape parameter k and scale parameter λ,

$$p(x,\lambda,k) = \frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k}, \quad x,\lambda,k > 0. \tag{3.122}$$

The variable x can be seen as a power of an exponentially distributed variable with intensity unity, $p(t) = e^{-t}$, for $t = \left(\frac{x}{\lambda}\right)^k$. It is known that the moment problem associated to the moments of x is determinate for $k \geq 1/2$ and indeterminate for $k < 1/2$ (Stoyanov, 1987, section 11.3 e.g.). A Stieltjes class in the latter regime is provided in the same reference, which after some algebraic manipulations reduces to

$$h(x) = \sin\left(c_k(x/\lambda)^k - k\pi\right), \quad c_k = \tan\left(k\pi\right), \tag{3.123}$$

The index D_S becomes

$$D_S = \int_0^\infty e^{-t}\left|\sin\left(c_k\, t - k\pi\right)\right| \, dt. \tag{3.124}$$

It is interesting to note that this integral can be performed analytically, for instance with the help of partial integration, with the result

$$D_S = \sin\left(k\, 2\pi\right)\frac{e^{-k\pi/c_k}}{1 - e^{-\pi/c_k}}. \tag{3.125}$$

We evaluated ϵ_N and the associated $F_{\leq N}$ numerically, using the convenient coordinates $\ln\lambda$ and $\ln k$. It is not difficult to see that ϵ_N is independent of λ. The full information matrix for these parameters can be evaluated analytically, with the result $\det F = k\pi^2/6$. The results for ϵ_N are shown in figure 3.2, that we present as for the lognormal as function of the reduced variance

$$\frac{\sigma}{\mu} = \left(\frac{\Gamma\left(1+2/k\right)}{\Gamma^2\left(1+1/k\right)} - 1\right)^{1/2}. \tag{3.126}$$

The dashed vertical line corresponding to $k = 1/2$ separates the two regimes where, on the left, ϵ is unity and D_S zero since the moment problem is determinate, and on the right, where the moment problem is indeterminate. The scale $\sigma/\mu = 1$ corresponds to $k = 1$. While the decay of ϵ_N is also very sharp in the indeterminate regime, very slow convergence of ϵ_N is seen to occur in the large reduced variance regime unlike for the lognormal distribution. This also in the region $1 < k < 2$, which corresponds to the phase where the Weibull distribution goes from a unimodal to a monotonically decreasing distribution, but ϵ still is unity, since the moment problem still is determinate. For instance, for $k = 1$, the following exact result can be gained with the same methods as exposed in this section,

$$\epsilon_N|_{k=1} = \left(\sum_{n=1}^{N-1}\frac{1}{n^2}\right)\frac{6}{\pi^2}. \tag{3.127}$$

Stretched exponential function Finally, we treat the case of the stretched exponential,

$$p(x,\lambda,k) = \frac{k}{\lambda\Gamma(1/k)}e^{-(x/\lambda)^k}, \quad x \geq 0. \tag{3.128}$$

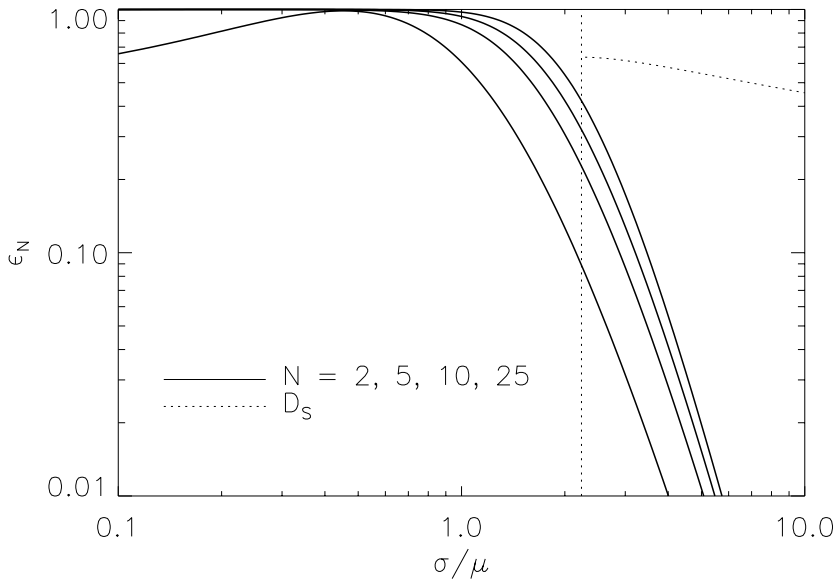

Fig. 3.2 — Solid lines : the cumulative efficiencies ϵ_N, for $N = 2, 5, 10, 25$, from bottom to top, as function of σ/μ for the Weibull distribution. Dashed : The dissimilarity index for the Stieltjes class given in equation (3.123). The vertical dashed line, corresponding to $k = 1/2$, separates the regimes of determinacy and indeterminacy of the associated moment problem.

Just as for the Weibull distribution, the moment problem associated to the moments of x is determinate for $k \geq 1/2$ and indeterminate for $k < 1/2$ (Stoyanov, 1987, section 11.4) A Stieltjes class is given by

$$h(x) = \sin\left(c_k (x/\lambda)^k\right), \quad x \geq 0,\ k > 0 \tag{3.129}$$

where c_k is as above given by $c_k = \tan(k\,\pi)$. Numerical evaluation of ϵ_N, and of the dissimilarity index is shown in fig 3.3. As for the Weibull distribution, these results are independent of the scale parameter λ. We can conclude that in none of the situations we investigated is D_S a good tracer of the importance of the indeterminacy of the moment problem for parameter inference.

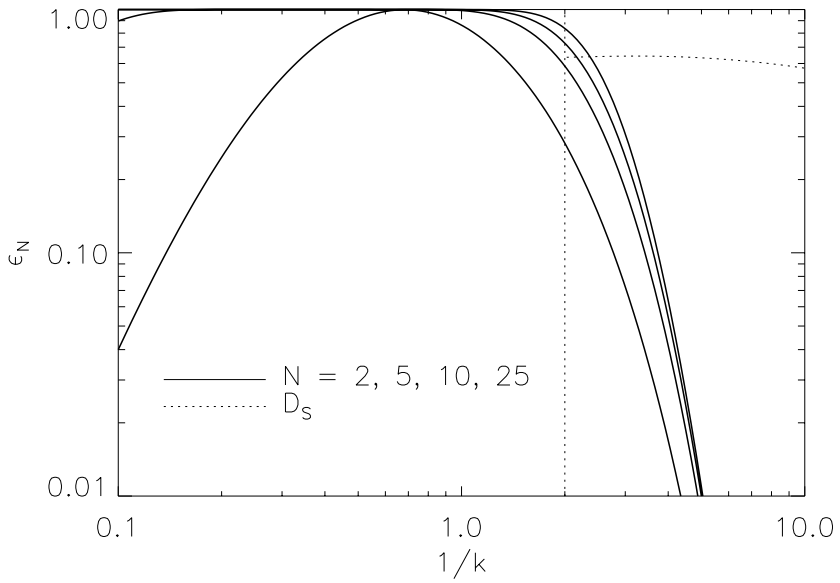

Fig. 3.3 — Solid lines : the cumulative efficiencies ϵ_N, for $N = 2, 5, 10, 25$, from bottom to top, as function of the inverse shape parameter $1/k$ of the stretched exponential. Dashed : The dissimilarity index for the Stieltjes class given in equation (3.129). The vertical dashed line at $k = 1/2$ separates the regimes of determinacy and indeterminacy of the associated moment problem.

3.4 Appendix

3.4.1 Derivation for the beta family

In the following we will need the first two derivatives of the logarithm of the gamma function. These are called the digamma ψ_0 and trigamma ψ_1 functions respectively (Abramowitz and Stegun, 1970, p. 258-260),

$$\begin{aligned}\psi_0(x) &= \frac{d}{dx}\ln\Gamma(x)\\ \psi_1(x) &= \frac{d^2}{dx^2}\ln\Gamma(x).\end{aligned} \tag{3.130}$$

The Fisher information matrix can be gained by noting that by differentiation under the integral sign, we have

$$F_{\alpha\beta} = \frac{\partial^2 \ln B(\alpha,\beta)}{\partial\alpha\partial\beta}. \tag{3.131}$$

Using the representation of the beta integral in terms of the Gamma function, we conclude that

$$F = \begin{pmatrix}\psi_1(\alpha) & 0\\ 0 & \psi_1(\beta)\end{pmatrix} - \psi_1(\alpha+\beta)\begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix}. \tag{3.132}$$

In order to obtain the information coefficients $s_n(\alpha)$ and $s_n(\beta)$, it is more convenient to start with $s_n(\beta)$. $s_n(\alpha)$ will then be gained effortlessly by looking at the symmetry of the problem.

From the definition of the beta integral, the moments of the beta distribution are given by

$$m_k = \frac{B(\alpha+k,\beta)}{B(\alpha,\beta)}. \tag{3.133}$$

The derivatives of the moments with respect to β are given by

$$\partial_\beta \ln m_k = -\psi_0(\alpha+\beta+k) + \psi_0(\alpha+\beta). \tag{3.134}$$

The orthogonal polynomials are the Jacobi polynomials $G_n^{(\alpha,\beta)}(x)$ (Abramowitz and Stegun, 1970, page 774). Instead of the parameters p, q used in Abramowitz and Stegun (1970), we stick to α and β, which are more appropriate for our purposes [4]. These polynomials are proportional to the Jacobi polynomials $P_n^{(\alpha-1,\beta-1)}(2x-1)$ (Szegö, 2003, chap. IV), orthogonal on the interval $(-1,1)$. Their matrix elements of G_n are given explicitly by

$$C_{nk} = \frac{\Gamma(\alpha+n)}{\Gamma(\alpha+\beta+2n-1)}(-1)^k\binom{n}{k}\frac{\Gamma(\alpha+\beta+k+n-1)}{\Gamma(\alpha+k)}. \tag{3.135}$$

[4] We have in the notation of Abramowitz and Stegun (1970) $q = \alpha, p = \alpha+\beta-1$

We first note the following relation,

$$C_{nk}m_k = \frac{\Gamma(\alpha+n)\Gamma(\beta)}{\Gamma(2n+\alpha+\beta-1)B(\alpha,\beta)}(-1)^{n-k}\binom{n}{k}(\alpha+\beta+k)^{(n-1)}, \tag{3.136}$$

where

$$(x)^{(n)} = x(x+1)\cdots(x+n-1) \tag{3.137}$$

is the rising factorial. Since $s_n(\beta)$ is given by

$$s_n(\beta) = \sum_{k=0}^{n} C_{nk}m_k\partial_\beta \ln m_k, \tag{3.138}$$

the evaluation of the following sum is necessary,

$$A_n(q) := \sum_{k=0}^{n}\binom{n}{k}(-1)^{n-k}(q+k)^{(n-1)}\left(\psi_0(q+k)-\psi_0(q)\right), \quad q := \alpha+\beta > 0. \tag{3.139}$$

The following paragraphs are dedicated to the lengthy but straightforward proof of the following result,

$$A_n(q) = \frac{(n-1)!}{q+n-1}. \tag{3.140}$$

The proof consists of a number of steps.

Step 1 : Using an always useful trick, we turn the rising factorial in the sum into a power and differentiate : from

$$(q+k)^{(n-1)} = (q+k)\cdots(q+k+n-2) = \left.\frac{d^{n-1}}{dt^{n-1}}\right|_{t=1} t^{k+q+n-2}, \tag{3.141}$$

follows

$$\begin{aligned} A_n(q) &= \left.\frac{d^{n-1}}{dt^{n-1}}\right|_{t=1} t^{q+n-2}\sum_{k=0}^{n}\binom{n}{k}(-1)^{n-k}t^k(\psi(q+k)-\psi(q)) \\ &=: \left.\frac{d^{n-1}}{dt^{n-1}}\right|_{t=1} t^{q+n-2}f_q(t). \end{aligned} \tag{3.142}$$

We then perform the derivatives using the Leibniz rule of derivation. Since

$$\left.\frac{d^{n-1-m}}{dt^{n-1-m}}\right|_{t=1} t^{q+n-2} = \frac{\Gamma(q+n-1)}{\Gamma(q+m)}, \tag{3.143}$$

one obtains

$$A_n(q) = \Gamma(q+n-1)\sum_{m=0}^{n-1}\binom{n-1}{m}\frac{f_q^{(m)}(t=1)}{\Gamma(q+m)}. \tag{3.144}$$

Step 2 To obtain $f_q^{(m)}(1)$, we first construct an integral representation of the function $f_q(t)$. From the relation (Andrews et al., 1999, theorem 1.2.7)

$$\psi_0(q+1) = \frac{1}{q} + \psi_0(q) \tag{3.145}$$

of the digamma function, we note that

$$\psi_0(q+k) - \psi_0(q) = \sum_{j=0}^{k-1} \frac{1}{q+j}. \tag{3.146}$$

Writing

$$\frac{1}{q+j} = \int_0^1 dx\, x^{q+j-1}, \tag{3.147}$$

the sum over j becomes a geometric series. We obtain therefore,

$$\psi_0(q+k) - \psi_0(q) = \int_0^1 dx\, x^{q-1} \frac{1-x^k}{1-x}. \tag{3.148}$$

We can now perform the sum over k in $f_q(t)$. We have indeed

$$\sum_{k=1}^{n} \binom{n}{k} (-1)^{n-k} t^k (1-x^k) = (t-1)^n + (-1)^n (1-tx)^n. \tag{3.149}$$

It follows immediately

$$f_q(t) = \int_0^1 dx\, x^{q-1} \frac{(t-1)^n - (-1)^n (1-tx)^n}{1-x}. \tag{3.150}$$

We only need derivatives of that function evaluated at $t = 1$. The order of each derivative is $< n$. For this reason the term $(t-1)^n$ in this expression actually plays no role. Each such derivative is thus a beta integral :

$$f_q^{(m)}(1) = (-1)^{n+m-1} n \cdots (n-m+1) B(q+m, n-m), \quad m < n. \tag{3.151}$$

Step 3 : To go further, we use the following property of the beta integral,

$$B(x,y) = \frac{(x+y)^{(n)}}{(y)^{(n)}} B(x, y+n), \tag{3.152}$$

which can be seen from its representation in terms of the gamma function, or from its integral representation (Andrews et al., 1999, page 5). Two applications of this rule leads to

$$B(q+m, n-m) = B(q,n) \frac{(q)^{(m)}}{m!} \frac{1}{\binom{n-1}{m}}. \tag{3.153}$$

By combining this relation with (3.151), and using

$$(q)^{(m)} = \frac{\Gamma(q+m)}{\Gamma(q)}, \tag{3.154}$$

we have shown

$$f^{(m)}(1) = (-1)^{n+m-1} B(q,n) \frac{\Gamma(q+m)}{\Gamma(q)} \binom{n}{m} \frac{1}{\binom{n-1}{m}}. \tag{3.155}$$

This form is very convenient, since many terms now cancel in equation (3.144). We obtain

$$A_n(q) = -\underbrace{B(q,n)\frac{\Gamma(n+q-1)}{\Gamma(q)}}_{\frac{(n-1)!}{n+q-1}} \sum_{m=0}^{n-1} \binom{n}{m} (-1)^{n-m}. \tag{3.156}$$

The sum over m would vanish would it run up to n. Its value is therefore minus the $m = n$ term, which is unity. It follows

$$A_n(q) = \frac{(n-1)!}{n+q-1}, \tag{3.157}$$

which was to be proved.

Getting $s_n(\beta)$ requires now only to keep track of the normalization. The normalization of the Jacobi polynomials G_n Abramowitz and Stegun (1970) is

$$h_n = \frac{n!}{B(\alpha,\beta)} \frac{\Gamma(\alpha+n)\Gamma(\alpha+\beta-1+n)\Gamma(\beta+n)}{(\alpha+\beta-1+2n)\Gamma^2(\alpha+\beta-1+2n)}. \tag{3.158}$$

(Note that there is a additional factor of $1/B$ with respect to Abramowitz and Stegun (1970) since there the measure is not normalized to unit integral). We have from equations (3.134), (3.136), (3.138) and (3.140) together with some algebra

$$s_n(\beta) = -\frac{1}{n} \left(\frac{(\alpha)^{(n)} n!}{(\beta)^{(n)} (\alpha+\beta)^{(n)}} \right)^{1/2} \left(\frac{2n+\alpha+\beta-1}{n+\alpha+\beta-1} \right)^{1/2}. \tag{3.159}$$

Tedious calculations are not needed to get $s_n(\alpha)$, but symmetry considerations are enough. The Jacobi polynomials obeys the symmetry relation,

$$P_n^{(\alpha-1,\beta-1)}(-x) = (-1)^n P_n^{(\beta-1,\alpha-1)}(x), \tag{3.160}$$

and therefore

$$G_n^{(\alpha,\beta)}(1-x) = (-1)^n G_n^{(\beta,\alpha)}(x). \tag{3.161}$$

It is then not difficult[5] to see that $s_n(\alpha)$ is proportional to $s_n(\beta)$ with α and β exchanged. We conclude

$$s_n(\alpha) = -\frac{(-1)^n}{n} \left(\frac{(\beta)^{(n)} n!}{(\alpha)^{(n)} (\alpha+\beta)^{(n)}} \right)^{1/2} \left(\frac{2n+\alpha+\beta-1}{n+\alpha+\beta-1} \right)^{1/2}. \tag{3.162}$$

[5]for instance from the representation (3.9)

3.4.2 General hierarchical systems, recursion relations.

This chapter was focussed on the hierarchy of N-point moments, with associated orthogonal system the orthogonal polynomials. One of course expects the approach of this chapter to extend in some way to any system of observables. It is the aim of this section to discuss briefly the case of other hierarchical systems.

Given a density function $p(x, \boldsymbol{\theta})$, where $x = (x_1, \cdots, x_d)$, consider a set of functions, (a *hierarchy* $[f]$),

$$[f] = f_0(x), f_1(x), f_2(x), \cdots \tag{3.163}$$

which can be finite or infinite. The restrictions are $f_0(x) = 1$, the functions to be linearly independent with respect to the density $p(x, \boldsymbol{\theta})$, as well as $\langle f_i f_j \rangle$ to be well defined.

The density provides us with a scalar product,

$$f \cdot g := \langle f(x) g(x) \rangle = \int dx\, p(x, \boldsymbol{\theta}) f(x) g(x). \tag{3.164}$$

Consider now the orthonormal system built from $[f]$, following the Gram-Schmidt orthogonalisation procedure. Explicitly, they can be built recursively from

$$P_n(x) = \frac{f_n(x) - \sum_{k=0}^{n} \langle f_n P_k \rangle P_k(x)}{\sqrt{\left\langle \left(f_n(x) - \sum_{k=0}^{n} \langle f_n P_k \rangle P_k(x)\right)^2 \right\rangle}} \tag{3.165}$$

By construction, we have indeed

$$\langle P_n(x) P_m(x) \rangle = \delta_{mn}. \tag{3.166}$$

P_n is a polynomial in the hierarchy members, in the sense that

$$P_n(x) = \sum_{k=0}^{n} C_{nk} f_k(x), \tag{3.167}$$

for some matrix elements C_{nk}, with $C_{nn} > 0$, and $C_{nk} = 0, \quad k > n$.

Very useful is the following identity between matrices of size $N + 1 \times N + 1$:

$$C^T C = [M_N]^{-1}, \quad [M_N]_{ij} = \langle f_i(x) f_j(x) \rangle, \quad i, j = 0, \cdots N. \tag{3.168}$$

Derivation Consider the expansion of f_k in terms of P_n :

$$f_k(x) = \sum_{n=0}^{k} B_{kn} P_n(x). \tag{3.169}$$

Such an expansion is always possible since the set of P_n and f_n span the same space by construction. It follows $B_{kn} = \langle f_k P_n \rangle$. Expanding in this relation P_n with the help of the matrix C, we obtain

$$f_k(x) = \sum_{n=0}^{k} \sum_{i,j=0}^{n} C_{ni} C_{nj} \langle f_i f_k \rangle f_j(x), \quad k = 0, 1, \cdots, \forall x \tag{3.170}$$

Relation (3.168) follows.

We are writing now the moments of the hierarchy as $m_k := \langle f_k \rangle$. The inverse of the covariance matrix $[\Sigma_N]_{ij} = \langle f_i f_j \rangle - m_i m_j$ between the members of $[h]$ is given by

$$\left[\Sigma_N^{-1}\right]_{ij} = \left[C^T C\right]_{ij} = \sum_{n=0}^{N} C_{ni} C_{nj}, \quad i, j = 1, \cdots, N \tag{3.171}$$

This is indeed the $N \times N$ lower right block of M_N^{-1}.

From these considerations follows : Define the (positive) Fisher information matrix within the first N hierarchy members, $[F_N]$ as

$$[F_{\leq N}]_{\alpha\beta} := \sum_{i,j=0}^{N} \frac{\partial m_i}{\partial \alpha} \left[M_N^{-1}\right]_{ij} \frac{\partial m_j}{\partial \beta}. \tag{3.172}$$

If the density is a probability density, then $\partial_\alpha m_0 = 0$, and therefore

$$[F_{\leq N}]_{\alpha\beta} = \sum_{i,j=1}^{N} \frac{\partial m_i}{\partial \alpha} \left[\Sigma_N^{-1}\right]_{ij} \frac{\partial m_j}{\partial \beta}. \tag{3.173}$$

Writing M_N^{-1} as $C^T C$ we can expand the matrix as

$$[F_{\leq N}]_{\alpha\beta} = \sum_{n=1}^{N} s_n(\alpha) s_n(\beta), \text{ with } s_n(\alpha) = \left\langle \frac{\partial \ln p}{\partial_\alpha} P_n \right\rangle = \sum_{k=0}^{n} C_{nk} \frac{\partial m_k}{\partial \alpha}. \tag{3.174}$$

The coefficients $s_n(\alpha)$ is therefore the information on α in f_n that was not contained already in the previous members of the hierarchy. The bits of information absent from the hierarchy are given by

$$\left\langle \left(\frac{\partial \ln p}{\partial \alpha} - \sum_{k=1}^{n} s_n(\alpha) P_n(x) \right) \left(\frac{\partial \ln p}{\partial \beta} - \sum_{k=1}^{n} s_n(\beta) P_n(x) \right) \right\rangle = F_{\alpha\beta} - [F_{\leq N}]_{\alpha\beta}\,. \tag{3.175}$$

This matrix on the right hand side of that relation is positive (it is a Gram matrix), and thus this relation states that the information within the hierarchy is always less than the total information. From this inequality follows thus the information inequality for any set of probes. Convergence to the total information occurs when

$$s_{\leq N}(x, \alpha) := \sum_{n=0}^{N} s_n(\alpha) P_n(x) \overset{N \to \infty}{\to} \frac{\partial \ln p(x, \boldsymbol{\theta})}{\partial \alpha}, \tag{3.176}$$

in the mean square error sense with respect to $p(x, \boldsymbol{\theta})$. This implies that the hierarchy is efficient precisely when $\partial_\alpha \ln p$ is sparse in $[f]$. In particular, maximal entropy distributions with the prescribed moments of $[f]$ have a finite non zero numbers of coefficients. Note that $s_{\leq N}(x, \alpha)$ is the best approximation of $\partial_\alpha \ln p$ with the given hierarchy up to N according to this mean square criterium.

Recursion relations Remember the triangular matrix B of size N defined above, $B_{nk} = \langle f_n P_k \rangle$. The following relations are easily seen to hold,

$$BB^T = M_N, \quad B = C^{-1}. \tag{3.177}$$

It follows that the matrix B is nothing else than the Cholesky decomposition of the moment matrix. Multiplying (3.165) with the score function and integrating, one obtains

$$s_n(\alpha) = \frac{1}{B_{nn}} \left(\frac{\partial m_n}{\partial \alpha} - \sum_{k=0}^{n} B_{nk}\, s_k(\alpha) \right). \tag{3.178}$$

This can provide a way to evaluate numerically these coefficient, after a Cholesky decomposition of the moment matrix. However, it is necessary to keep in mind that large moment matrices are infamously known for being generically badly conditioned (Tyrtyshnikov, 1994).

Part II

Applications for cosmology

4 On the combination of shears, magnification and flexion fields

At the heart of weak lensing as a cosmological probe lies the convergence field, the weighted projection along the line of sight of of the density fluctuation field, with weights sensitive to the geometry of the Universe. In this chapter we use the duality of the Shannon entropy and Fisher information introduced in chapter 2 to discuss quantitatively the combination of different probes of the convergence field, both for mass reconstruction as well as for cosmological purposes. The probes we include are the galaxy shears, the magnification as well as the flexion fields, i.e. all modifications to the galaxy images up to second order. Section 4.2 describes the observables and the approach, and section 4.3 the results.

We find that flexion alone outperforms the well established shears on the arcsecond scale, making flexion well suited for mass reconstruction on small scales. At the same time, it complements powerfully the shears on the scale of the arcminute. We find size information to carry some modest, scale independent amount of information. Besides, the results of this chapter show how standard cosmological Fisher matrix analysis based on Gaussian statistics can incorporate these other probes in the most simple way. From (4.24) follows namely that the inclusion of all the two-point correlations of these additional weak lensing probes can be accounted for by adapting the noise term.

The text in this section follows the second part of Carron et al. (2011).

4.1 Introduction

Gravitational lensing, which can be used to measure the distribution of mass along the line of sight, has been recognized as powerful probe of the dark components of the Universe

(Bartelmann and Schneider, 2001; Munshi et al., 2006; Refregier, 2003; Schneider et al., 1992, 2006) since it is sensitive to both the geometry of the Universe, and to the growth of structure. Weak lensing data is typically used in two ways. The first, which is deployed for cosmological parameter fitting, relies on measuring the correlated distortions in galaxy images (Albrecht et al., 2006). The second approach uses each galaxy to make a noisy measurement of the lensing signal at that position. These point estimates are then used to reconstruct the dark matter density distribution (e.g. Kaiser and Squires, 1993; Seitz and Schneider, 2001). Most of the measurements of weak lensing to date have focused on the shearing that galaxy images experience. However, gravitational lensing causes a number of other distortions of galaxy images. These include change in size, which is related to the magnification, and higher order image distortions known as the flexion (Bacon et al., 2006). A number of techniques have been developed for measuring these higher order images distortions, such as HOLICS (Okura et al., 2007) and shapelets methods (Massey et al., 2007). Since all of the image distortions originate from the same source, the lensing potential field, the information content of any two lensing measurements must be degenerate. At the same time, since each method has different systematics and specific noise properties, combining multiple measurement may bring substantial benefits. Some recent works have looked at the impact of combining shear and flexion measurements for mass reconstruction (Er et al., 2010; Pires and Amara, 2010; Velander et al., 2010), as well as the benefits for breaking multiplicative bias of including galaxy size measurements (Vallinotto et al., 2010).

We will focus on tracers of galaxy image distortions up to second order in the distortions. These include, to first order, the change in apparent size of the galaxies, due to the magnification source by the convergence field, the two components of the shear, and, to second order, the four components of the two flexion fields. We will limit ourselves to the case where the noise contaminating each probe can be effectively treated as independent of the model parameters. The common point of the probes cited above is that they all trace, in some noisy way, the same central field on a discrete set of points, represented by the positions of the galaxies on which the tracers are measured. The framework we presented in chapter 2 is ideal to deal with the special structure of the situation. We will first show how we can understand the total information content of such degenerate probes of a central field in a general situation, and then make quantitative evaluations at the two-point level.

4.2 Linear tracers

The situation that we consider here is one where a broad number of observables are linked in some way to a central field. We will limit ourselves to the case where the noise contaminating each probe can be effectively treated as independent of the model parameters. Imagine one plans to perform different measurements, of observables that are all linked in some way or another to the same central field, which is the actually interesting quantity from the point of view of the analyst. The most straightforward example of such a situa-

tion, occurs when the two point correlations of the central field, say the the power spectrum of the convergence in weak lensing, is predicted from theory, while we try to extract it looking at the correlations of the two ellipticity components of galaxies. In this case, predictive power of the power spectrum of the convergence turns into predictive power of (parts of) the three correlation functions there are between the two ellipticity components.

The predictive power of some observable O_c of a central field (for instance its power spectrum at some mode) translates into an array of constraints $O_i,\ i = 1, \cdots, n$ in the noisy probes, that we could try and extract and confront to theory :

$$O_i(\boldsymbol{\theta}) = f_i(O_c(\boldsymbol{\theta})), \quad i = 1, \cdots, n \tag{4.1}$$

for some functions f_i.
For the purpose of this work, the case of functions linear with respect to O_c is generic enough, i.e. we will consider that

$$\frac{\partial^2 f_i}{\partial O_c^2} = 0, \quad i = 1, \cdots, n. \tag{4.2}$$

The entropy S of the data is a function of the n constraints $\mathbf{O}$. It is however fundamentally a function of O_c since it does enter all of these observables. It is therefore very natural to associate a potential λ_c to O_c, although it is not itself a constraint on the probability density function. In analogy with

$$\lambda_i = \frac{\partial S}{\partial O_i}, \quad i = 1, \cdots, n \tag{4.3}$$

we define

$$\lambda_c := \frac{dS}{dO_c}(O_1, \cdots, O_m), \tag{4.4}$$

with the result, given by application of the chain rule, of

$$\lambda_c = \boldsymbol{\lambda} \cdot \frac{\partial \mathbf{f}}{\partial O_c}. \tag{4.5}$$

On the other hand, the impact of a model parameter on each observables can be similarly written in terms of the central observable O_c,

$$\frac{\partial \mathbf{O}}{\partial \alpha} = \frac{\partial O_c}{\partial \alpha} \frac{\partial \mathbf{f}}{\partial O_c}. \tag{4.6}$$

It follows directly from these relations (4.5) and (4.6), and the linearity of f_i, that the joint information in the full set of constraints $\mathbf{O}$, given in equation (2.43) as a sum over all n constraints, reduces to a formally identical expression with the only difference that only O_c enters :

$$F^X_{\alpha\beta} = \frac{\partial \mathbf{O}}{\partial \alpha} \cdot \frac{\partial \boldsymbol{\lambda}}{\partial \beta} = \frac{\partial \lambda_c}{\partial \alpha} \frac{\partial O_c}{\partial \beta}, \tag{4.7}$$

which can also be written in the form analog to (2.44),

$$F^X_{\alpha\beta} = -\frac{\partial O_c}{\partial \alpha} \frac{d^2 S}{dO_c^2} \frac{\partial O_c}{\partial \beta}. \tag{4.8}$$

This last equation shows that all the effect of combining this set of constraints have been absorbed into the second total derivative of the entropy. This second total derivative is the total amount of information there is on the central quantity O_c in the data. Indeed, taking as a special case of model parameter to the central quantity itself, i.e.

$$\alpha = \beta = O_c, \tag{4.9}$$

one obtains now that the full amount of information in X on O_c is

$$F^X_{O_c O_c} = -\frac{d^2 S}{dO_c^2}\left(O_1, \cdots, O_n\right) \equiv \frac{1}{\sigma^2_{\text{eff}}}. \tag{4.10}$$

A simple application of the Cramer Rao inequality presented in equation (2.10) shows that this effective variance σ_{eff} is the lower bound to an unbiased reconstruction of the central observable from the noisy probes.

These considerations on the effect of probe combination in the case of a single central field observable O_c generalize easily to the case where there are many, $(O_c^1, \cdots, O_c^m)$. In this case, each central field quantity leads to an array of constraints in the form of equation (4.1), it is simple to show that the amount of Fisher information can again be written in terms of the information associated to the central field, with an effective covariance matrix between the $O_c's$. The result is

$$F^X_{\alpha\beta} = -\sum_{i,j=1}^{m} \frac{\partial O_c^i}{\partial \alpha} \frac{d^2 S}{dO_c^i O_c^j} \frac{\partial O_c^j}{\partial \beta}. \tag{4.11}$$

All the effects of probe combination are thus encompassed in an effective covariance matrix Σ_{eff} of the central field observables,

$$-\frac{d^2 S}{dO_c^i O_c^j} \equiv \left[\Sigma^{-1}_{\text{eff}}\right]_{ij}. \tag{4.12}$$

Again, an application of the Cramer Rao inequality, in the multi-dimensional case, shows that this effective covariance matrix is the best achievable unbiased joint reconstruction of $(O_c^1, \cdots, O_c^m)$.

We now explore further the case of linear probes of homogeneous Gaussian fields, which is cosmologically relevant and can be solved analytically to full extent. We will focus on zero mean fields, for which according to our previous section the entropy can be written in terms of the spectral matrices, up to a constant,

$$S = \frac{1}{2}\sum_{\mathbf{k}} \ln\det P(\mathbf{k}). \tag{4.13}$$

4.2.1 Linear tracers at the two-point level

The standard instance of a linear tracer ϕ_i of some central field κ in weak lensing is provided by a relation in Fourier space of the form

$$\tilde{\phi}_i(\mathbf{k}) = v_i \tilde{\kappa}(\mathbf{k}) + \tilde{\epsilon}_i(\mathbf{k}), \tag{4.14}$$

for some noise term $\tilde{\epsilon}_i$, uncorrelated with κ, and coefficient v_i. Typically, if one observes a tracer of the derivative of the field κ, then the vector $\mathbf{v}$ would be proportional to $-i\mathbf{k}$. We are ignoring here any observational effect, such as incomplete sky coverage, that would require corrections to this relation. It is clear from this relation that the spectral matrices of this family take a special form of equation (4.1): defining the spectrum of the κ field by P^κ, we obtain by putting this relation (4.14) into (2.59), that the spectral matrices can be written at each mode in the form

$$P = P^\kappa \mathbf{v}\mathbf{v}^\dagger + N, \tag{4.15}$$

where $\mathbf{v}^\dagger$ is the hermitian conjugate of $\mathbf{v} = (v_1, \cdots, v_n)$. The matrix N is the spectrum of the noise components ϵ,

$$N_{ij}(\mathbf{k}) = \frac{1}{V} \left\langle \tilde{\epsilon}_i(\mathbf{k}) \tilde{\epsilon}^*(\mathbf{k}) \right\rangle . \tag{4.16}$$

Our subsequent results hold for any family of tracers that obey this relation. While the special case in (4.14) enter this category, this need not be the only instance. All the weak lensing observables we deal with in this work will satisfy equation (4.15).

Both the n-dimensional vector $\mathbf{v}$ and the noise matrix N can depend on the wave vector $\mathbf{k}$, but they are independent from the model parameters. The matrix N of dimension $n \times n$ is the noise component of the spectra of the fields, typically built from two parts. The first is due to the discrete nature of the fields, since such data consist in quantities measured where galaxies sits, and the second to the intrinsic dispersion of the measured values.

4.2.2 Joint entropy and information content

Information on the model parameters enters through P^κ only. To evaluate the full information content, we need only evaluate eq. (4.13) with the spectral matrix given in (4.15), keeping in mind the result from last section, that we need only the total derivative with respect to P^κ. In other words, any additive terms in the expression of the entropy that are independent of P^κ can be discarded.

This determinant can be evaluated immediately. Defining for each mode the real positive number N_{eff} through

$$\frac{1}{N_{\text{eff}}} \equiv \mathbf{v}^\dagger N^{-1} \mathbf{v}, \tag{4.17}$$

which can be seen as an effective noise term, a simple[1] calculation shows that the joint entropy (4.13) is equivalent to the following, where the n dimensional determinant has disappeared,

$$S = \frac{1}{2}\sum_{\mathbf{k}} \ln\left(P^{\kappa}(\mathbf{k}) + N_{\rm eff}(\mathbf{k})\right). \tag{4.19}$$

Comparison with equation (4.13) shows that we have with this equation (4.19) the entropy of the field κ itself, where all the effects of the joint observation of this n fields have been absorbed into the effective noise term $N_{\rm eff}$, that contaminates its spectrum. It means that the full combined information in the n probes of the field κ is equivalent to the information in κ, observed with spectral noise $N_{\rm eff}$.

Our result (4.10) applied to (4.19) puts bounds on reconstruction of the field κ out of the observed samples, which can be at best reconstructed with a contaminating noise term of $N_{\rm eff}$ in its spectrum, whose best unbiased reconstruction is given by

$$2\left(P^{\kappa}(\mathbf{k}) + N_{\rm eff}(\mathbf{k})\right)^2. \tag{4.20}$$

Since the effect of combining these probes at a single mode is only to change the model independent noise term, the parameter correlations and degenaracies as approximated by the Fisher information matrix stay unchanged, whatever the number of such probes is. We have namely from (4.19) that at a given mode $\mathbf{k}$, the Fisher information matrix reads

$$F^X_{\alpha\beta} = \frac{1}{2}\frac{\partial \ln \tilde{P}^{\kappa}(\mathbf{k})}{\partial\alpha}\frac{\partial \ln \tilde{P}^{\kappa}(\mathbf{k})}{\partial\beta}, \tag{4.21}$$

with

$$\tilde{P}^{\kappa}(\mathbf{k}) = P^{\kappa}(\mathbf{k}) + N_{\rm eff}(\mathbf{k}). \tag{4.22}$$

From the point of view of the Fisher information, it makes formally no difference to extract the full set of $n(n-1)/2$ independent elements of each spectral matrices, or reconstruct the field κ and extract its spectrum. They carry indeed the same amount of Fisher information.

These results still hold when other fields are present in the analysis, which are correlated with the field κ. To make this statement rigorous, consider in the analysis on top of our n samples of the form (4.14) of κ, another homogeneous field θ, with spectrum $P^{\theta}(\mathbf{k})$, and cross spectrum to κ given by $P^{\theta\kappa}(\mathbf{k})$ The full spectral matrices are in this case

$$P(\mathbf{k}) = \begin{pmatrix} P^{\kappa}(\mathbf{k})\mathbf{v}\mathbf{v}^T + N & P(\mathbf{k})^{\kappa\theta}\mathbf{v} \\ P^{\theta\kappa}\mathbf{v}^T & P^{\theta}(\mathbf{k}) \end{pmatrix}. \tag{4.23}$$

Again, the determinant of this matrix can be reduced to a determinant of lower dimension, leading to the equivalent entropy

$$S = \mathrm{cst} + \frac{1}{2}\ln\det\begin{pmatrix} P^{\psi}(\mathbf{k}) + N_{\rm eff} & P^{\kappa\theta}(\mathbf{k}) \\ P^{\theta\kappa}(\mathbf{k}) & P^{\theta}(\mathbf{k}) \end{pmatrix}. \tag{4.24}$$

[1] We have namely for any invertible matrix A and vectors $\mathbf{u}, \mathbf{v}$ the matrix determinant lemma,

$$\det\left(A + \mathbf{u}\mathbf{v}^T\right) = \det(A)\left(1 + \mathbf{v}^T A^{-1}\mathbf{u}\right). \tag{4.18}$$

It shows that the the full set of $n+1$ fields can be reduced without loss to two fields, κ and θ, with the effective noise N_{eff} contaminating the spectrum of κ.

Note that the derivation of our results do not refer to any hypothetical estimators, but came naturally out of the expression of the entropy.

4.2.3 Weak lensing probes

We now seek a quantitative evaluation of the full joint information content of the weak lensing probes in galaxy surveys, up to second order in the image distortions of galaxies. The data X consists of a set of fields, which are discrete point fields, which take values where galaxies sits. We work in the two-dimensional flat sky limit, using the more standard notation $\mathbf{l}$ for the wave vector, and decompose it in modulus and polar coordinate as

$$\boldsymbol{l} = l \begin{pmatrix} \cos\varphi_l \\ \sin\varphi_l \end{pmatrix} \tag{4.25}$$

We will throughout assume that the intrinsic values of each probe are pairwise uncorrelated, as commonly done. Also, we will assume that the set of points on which the relevant quantities are measured show low enough clustering so that corrections to the spectra due to intrinsic clustering can be neglected. This is however not a limitation of our approach, since corrections to the above assumptions, such as the introduction of some level of intrinsic alignment, can be accommodated for by introducing appropriate terms in the noise matrices $N(\mathbf{k})$ in (4.17). As a central field to which all our point fields relates, we take for convenience the isotropic convergence field κ, with spectrum

$$C^{\kappa}(\boldsymbol{l}) = C^{\kappa}(l). \tag{4.26}$$

In the case of pairwise uncorrelated intrinsic values that we are following, we see easily from (4.17) that by combining any number of such probes the effective noise is reduced at a given mode according to

$$\frac{1}{N_{\text{eff}}^{\text{tot}}} = \sum_i \frac{1}{N_{\text{eff}}^{i}}. \tag{4.27}$$

We therefore only need to evaluate the effective noise for each probe separately, while their combination follows (4.27). To this aim, the evaluation of the spectral matrices (4.15), giving us N_{eff}, is necessary. The calculations for this are presented in appendix 4.4 and we use the final results in this section.

First order, distortion matrix

To first order, the distortion induced by weak lensing on a galaxy image is described by the distortion matrix that contains the shear, γ, and convergence, κ, which come from the

second derivatives of the lensing potential field ψ, (e.g. Schneider et al., 2006)

$$\begin{pmatrix} \kappa+\gamma_1 & \gamma_2 \\ \gamma_2 & \kappa-\gamma_1 \end{pmatrix} = \psi_{,ij}. \tag{4.28}$$

The shear components read

$$\gamma_1 = \frac{1}{2}\left(\psi_{,11} - \psi_{,22}\right), \quad \gamma_2 = \psi_{,12} \tag{4.29}$$

and we assume they are measured from the apparent ellipticities of the galaxies, with identical intrinsic dispersion σ_γ^2. Denoting with $\bar{n}_\gamma$ the number density of galaxies for which ellipticity measurements are available, the effective noise is

$$N_{\rm eff}^{\gamma} = \frac{\sigma_\gamma^2}{\bar{n}_\gamma}. \tag{4.30}$$

The information content of the two observed ellipticity fields is thus exactly the same as the one of the convergence field, with a mode independent noise term as above.

To reach for the κ component of the distortion matrix, we imagine we have measurements of their angular size $s_{\rm obs}$, with intrinsic dispersion σ_s^2. The intrinsic sizes of the galaxies $s_{\rm int}$ gets transformed through weak lensing according to

$$s_{\rm obs} = s_{\rm int}(1 + \alpha_s \kappa). \tag{4.31}$$

The coefficient α_s, is equal to unity in pure weak lensing theory, but we allow it to take other values, since in a realistic situation, other effects such as magnification bias effectively enter this coefficient (see e.g. Vallinotto et al. (2010)). Under our assumption that the correlation of the fluctuations in intrinsic sizes can themselves be neglected, the effective noise reduces to

$$N_{\rm eff}^{s} = \frac{1}{\alpha_s^2}\left(\frac{\sigma_s}{\bar{s}_{\rm int}}\right)^2 \frac{1}{\bar{n}_s}. \tag{4.32}$$

This combination of α_s with the dispersion parameters $\bar{s}$ and σ_s becomes the only relevant parameter in our case, and not the value of each of them.

Second order, flexion

To second order, the distortions caused by lensing on the galaxies images are given by third order derivatives of the lensing potential. These are conveniently described by the spin 1 and spin 3 flexion components $\mathcal{F}$ and $\mathcal{G}$, which in the notation of (Schneider and Er, 2008) read

$$\begin{aligned} \mathcal{F} &= \frac{1}{2}\begin{pmatrix} \psi_{,111} + \psi_{,122} \\ \psi_{,112} + \psi_{,222} \end{pmatrix} \\ \mathcal{G} &= \frac{1}{2}\begin{pmatrix} \psi_{,111} - 3\psi_{,122} \\ 3\psi_{,112} - \psi_{,222} \end{pmatrix}, \end{aligned} \tag{4.33}$$

Tab. 4.1 — Dispersion parameters used in figure 4.1.

σ_γ	$\sigma_\mathcal{F}$ asec^{-1}	$\sigma_\mathcal{G}$ asec^{-1}	$\frac{1}{\alpha_s}\frac{\sigma_s}{\bar{s}}$
0.25	0.04	0.04	0.9

and are extracted from measurements with intrinsic dispersion $\sigma^2_\mathcal{F}$ and $\sigma^2_\mathcal{G}$. The effective noise is this time mode-dependent,

$$\frac{1}{N_{\text{eff}}^{\mathcal{FG}}} = l^2 \left(\frac{\bar{n}_F}{\sigma_F^2} + \frac{\bar{n}_G}{\sigma_G^2} \right). \tag{4.34}$$

4.3 Results

Figure 4.1 shows the ratio of the effective noise to the noise present considering the shear fields only, assuming the same number densities of galaxies for each probe, and the values for the intrinsic dispersion stated in table 4.1. The conversion multipole l (upper x-axis) to angular scale θ (lower x-axis) follows $\theta = \pi/(l + 1/2)$. We have adopted for the size dispersion parameters the numbers from (Vallinotto et al., 2010), who evaluated this number for the DES survey conditions (The Dark Energy Survey Collaboration, 2005). We refer to the discussion in (Pires and Amara, 2010) for our choice of flexion dispersion parameters. The curves on this figure are ratios independent of the galaxy number density. They are redshift independent as well, only to the extent that the dispersion in intrinsic values can be treated as such. We can draw two main conclusions from figure 4.1. First, flexion information beings to play role only at the smallest scales, on the arcsecond scales, where it takes over and becomes the most interesting probe. On the scale of 1 amin, it can bring substantial improvement over shear only analysis, but only in combination with the shears, and not on its own. This is in good agreement with the comparative analysis of the power of the flexion $\mathcal{F}$ field and shears fields for mass reconstruction done in Pires and Amara (2010), restricted to direct inversion methods. Second, the inclusion of size of galaxies into the analysis provides a density independent, scale independent, improvement factor of

$$\frac{N_{\text{eff}}^{\gamma}}{N_{\text{eff}}^{\gamma+s}} = 1 + \left(\frac{\sigma_\gamma \bar{s}}{\sigma_s \alpha_s} \right)^2, \tag{4.35}$$

which is close to a 10% improvement for the quoted numbers. Of course, the precise value depends on the dispersion parameters of the population considered.

For the purpose of measuring cosmological parameters rather than mass reconstruction, more interesting are the actual values of the Fisher information matrices. Since with any combination of such probes, these matrices are proportional to each other at a single mode,

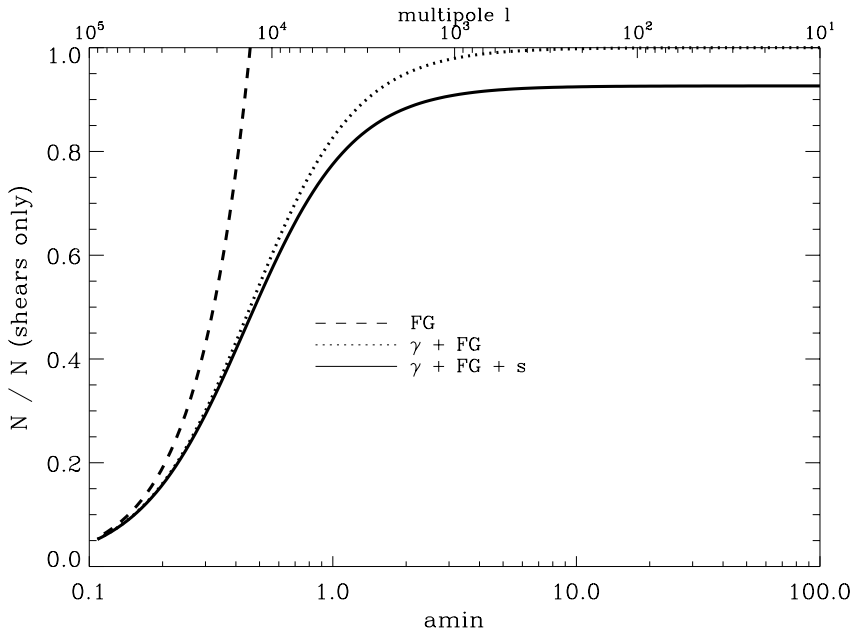

Fig. 4.1 — The ratio of the effective noise to the level of noise considering the shears only, as function of angular scale. The dashed line considers the flexion fields alone. The dotted line shows the combination of the flexion fields with the shear fields, and the solid line all these weak lensing probes combined. No correlations between the intrinsic values for each pair of probes have been considered.

it makes sense to define the efficiency parameter of the probe i through

$$\epsilon_i(l) := \frac{C^\kappa(l)}{C^\kappa(l) + N^i_{\rm eff}(l)}, \tag{4.36}$$

which is a measure of what fraction of the information contained in the convergence field is effectively captured by that probe. The information in the convergence field is, at a given mode l, counting the multiplicity of the mode,

$$F^\kappa_{\alpha\beta} = \frac{1}{2}(2l+1)\frac{\partial \ln C^\kappa(l)}{\partial\alpha}\frac{\partial \ln C^\kappa(l)}{\partial\beta}, \tag{4.37}$$

and we have indeed that the total Fisher information in the observed fields is

$$F^X_{\alpha\beta} = \sum_l F^\kappa_{\alpha\beta}(l)\epsilon_i^2(l). \tag{4.38}$$

Therefore, according to the interpretation of the Fisher matrix approximating the expected constraints on the model parameters, the factor $\epsilon(l)$ is precisely equal to the factor of degradation in the constraints one would be able to put on any a parameter, with respect to the case of perfect knowledge of the convergence field at this mode. It is not the purpose of this work to perform a very detailed study on the behavior of the efficiency

parameter for some specific survey and the subsequent statistical gain, but its qualitative behavior is easy to see. This parameter is essentially unity in the high signal to noise regime, while it is the inverse effective noise whenever the intrinsic dispersion dominates the observed spectrum. Since information on cosmological parameters is beaten down by cosmic variance in the former case, the latter dominates the constraints. We can therefore expect from our above discussion the size information to tighten by a few percent constraints on any cosmological parameter. On the other hand, while flexion becomes ideal for mass reconstruction purposes on small scales, it will be able to help inference on cosmological parameters only if the challenge of very accurate theoretical predictions on the convergence power spectrum for multipoles substantially larger than 1000 will be met.

To make these expectations more concrete, we evaluated the improvement in information on cosmological parameters performing a lensing Fisher matrix calculation for a wide, EUCLID-like survey, in a tomographic setting. For a data vector consisting of n probes of the convergence field κ_i in each redshift bin i, $i = 1, \cdots N$, it is simple to see following our previous argument, that the Fisher information reduces to

$$F_{\alpha\beta} = \frac{1}{2} \sum_l (2l+1) \, \mathrm{Tr} \, C^{-1} \frac{\partial C}{\partial \alpha} C^{-1} \frac{\partial C}{\partial \beta}, \tag{4.39}$$

where the C matrix is given by

$$C_{ij} = C^{\kappa_i \kappa_j}(l) + \delta_{ij} N^i_{\mathrm{eff}}(l), \quad i,j = 1, N \tag{4.40}$$

with N^i_{eff} given by (4.17). The only difference between standard implementations of Fisher matrices for lensing, such as the lensing part of Hu and Jain (2004), being thus the form of the noise component. we evaluated these matrices respectively for

$$N^i_{\mathrm{eff}} = \frac{\sigma^2_\gamma}{\bar{n}^i} = N^{\gamma,i}_{\mathrm{eff}}, \tag{4.41}$$

which is the precise form of the Fisher matrix for shear analysis, for

$$\frac{1}{N^i_{\mathrm{eff}}} = \frac{1}{N^{\gamma,i}_{\mathrm{eff}}} + \frac{1}{N^{s,i}_{\mathrm{eff}}} \tag{4.42}$$

which account for size information, and

$$\frac{1}{N^i_{\mathrm{eff}}(l)} = \frac{1}{N^{\gamma,i}_{\mathrm{eff}}} + \frac{1}{N^{s,i}_{\mathrm{eff}}} + \frac{1}{N^{\mathcal{FG},i}_{\mathrm{eff}}(l)}, \tag{4.43}$$

which accounts for the flexion fields as well. We note that in terms of observables, these small modifications incorporate in its entirety the full set of all possible correlations between the fields considered. The values of the dispersion parameters involved in these formulae are the same as in table (4.1). Our fiducial model is a flat ΛCDM universe, with parameters $\Omega_\Lambda = 0.7$, $\Omega_b = 0.045$, $\Omega_m = 0.3, h = 0.7$, power spectrum parameters $\sigma_8 = 0.8, n = 1$, and Chevallier- Polarski-Linder parametrisation (Chevallier and Polarski, 2001; Linder, 2003) of the dark energy equation of state implemented as $\omega_0 = -1, w_a = 0$. The distribution of galaxies as function of redshift needed both for the calculation of the

Tab. 4.2 — Ratio of the marginalised constraints $\sigma^2/\sigma^2_{\text{shear only}}$, for $l_{\max} = 10^4$. This first line considers the inclusion of the size information in the analysis, while the second the size as well as the flexion fields $\mathcal{F}$ and $\mathcal{G}$.

	Ω_Λ	Ω_b	Ω_m	h	n	σ_8	w_0	w_a
.	0.90	0.96	0.90	0.95	0.95	0.90	0.90	0.90
	0.88	0.96	0.89	0.95	0.93	0.88	0.88	0.88

spectra and to obtain the galaxy densities in each bin was generated using the cosmological package iCosmo (Refregier et al., 2008), in a way described in (Amara and Réfrégier, 2007). We adopted EUCLID-like parameters of 10 redshift bins, a median redshift of 1, a galaxy angular density of $40/\text{amin}^2$, and photometric redshift errors of $0.03(1+z)$.

In figure 4.2, we show the improvement in the dark energy Figure of Merit (FOM), defined as the square root of the determinant of the submatrix (ω_0, ω_a) of the Fisher matrix inverse $F^{-1}_{\alpha\beta}$ (α and β running over the set of eight parameters as described above), as function of the maximal angular mode $l_{\max}$ considered, while $l_{\min}$ being always taken to be 10. In perfect agreement with our discussion above, including size information (solid line) increases the FOM steadily until it saturates at a 10% improvement when constraints on the dark energy parameters are dominated by the low signal to noise regime. Also, flexion becomes only useful in the deep non-linear regime, where however theoretical understanding of the shape of the spectra still leaves a lot to be desired.

These results are found to be very insensitive to the survey parameters, for a fixed α_s. There are also only weakly model parameter independent, as illustrated in table 4.2, which shows the corresponding improvement in Fisher constraints,

$$\frac{\sigma^2}{\sigma^2_{\text{shear only}}} = \frac{F^{-1}_{\alpha\alpha}}{F^{-1}_{\alpha\alpha,\text{shear only}}}, \tag{4.44}$$

at the saturation scale $l_{\max} = 10^4$. These results are also essentially unchanged using either standard implementations of the halo model (Cooray and Sheth, 2002, for a review) or the the HALOFIT (Smith et al., 2003) non linear power spectrum.

4.4 Appendix

The data consists in a set of numbers, at each position where a galaxy sit and a measurement was done. We use the handy notation in terms of Dirac delta function,

$$\phi(\mathbf{x}) = \sum_i \epsilon_i \delta^D(\mathbf{x} - \mathbf{x}_i), \tag{4.45}$$

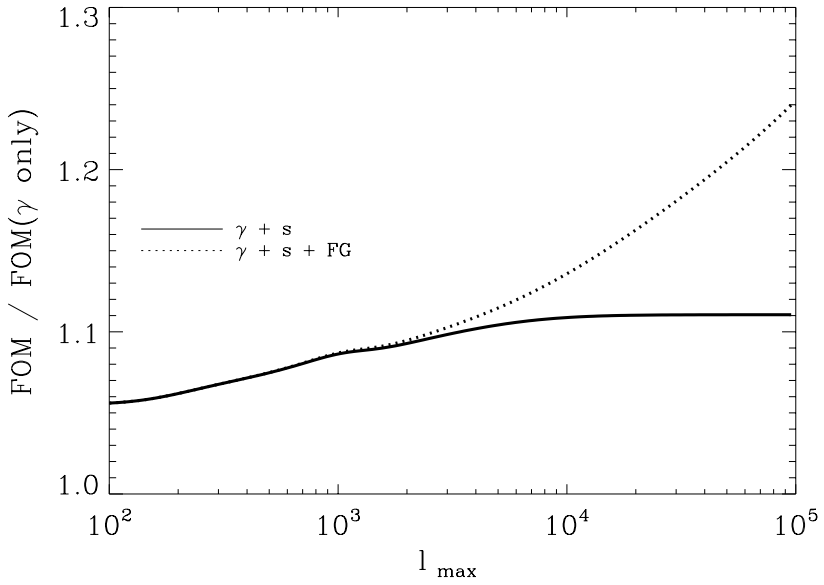

Fig. 4.2 — The improvement of the dark energy FOM including size information (solid), as well as flexion $\mathcal{F}$ and $\mathcal{G}$ information (dotted), over the shear only analysis, as function of the maximal angular multipole included in the analysis.

where the sum runs over the positions $\mathbf{x}_i$ for which ϵ is measured. To obtain the spectral matrices, we need the Fourier transform of the field, which reads in our case

$$\tilde{\phi}(\boldsymbol{l}) = \sum_i \epsilon_i \exp\left(-i\boldsymbol{l}\cdot\mathbf{x}_i\right). \tag{4.46}$$

In this work, we assume that the set of points shows negligible clustering, so that the probability density function for the joint occurrence of a particular set of galaxy positions is uniform.
We decompose in the following the wave vector $\mathbf{k}$ on the flat sky in terms of its modulus and polar angle as

$$\boldsymbol{l} = l\begin{pmatrix}\cos\phi_l\\ \sin\phi_l\end{pmatrix}. \tag{4.47}$$

Ellipticities When the two ellipticity components are measured, we have two such fields ϕ_1, ϕ_2 at our disposal. For instance, the field describing the first component becomes

$$\tilde{\phi}_1(\boldsymbol{l}) = \sum_i \epsilon_i^1 \exp\left(-i\boldsymbol{l}\cdot\mathbf{x}_i\right). \tag{4.48}$$

We assume that the measured ellipticities trace the shear fields, in the sense that the measured components are built out of the shear at that position plus some value unrelated

to it,

$$\begin{aligned} \epsilon_i^1 &= \gamma_1(\mathbf{x}_i) + \epsilon^1_{\mathrm{int},\,i} \\ \epsilon_i^2 &= \gamma_2(\mathbf{x}_i) + \epsilon^2_{\mathrm{int},\,i}. \end{aligned} \tag{4.49}$$

The vector $\mathbf{v}$ relating the spectral matrices of the ellipticities and the convergence is then obtained by plugging (4.48) with the above relations (4.49) in its definition (4.15), and using the relation between shears and convergence in equation (4.28). The result is

$$\mathbf{v} = \bar{n}_\gamma \begin{pmatrix} \cos 2\phi_l \\ \sin 2\phi_l \end{pmatrix}. \tag{4.50}$$

where $\bar{n}_\gamma$ is the number density of galaxies for which ellipticity measurements are available. Under our assumptions of uncorrelated intrinsic ellipticities, with dispersions of equal magnitude σ_γ^2 for the two components, the noise matrix N becomes

$$N = \bar{n}_\gamma \begin{pmatrix} \sigma_\gamma^2 & 0 \\ 0 & \sigma_\gamma^2 \end{pmatrix}. \tag{4.51}$$

The effective noise, given in equation (4.17) is readily computed

$$N^\gamma_{\mathrm{eff}} = \frac{\sigma_\gamma^2}{\bar{n}_\gamma}. \tag{4.52}$$

Sizes As noted in the main text, the apparent sizes of galaxies are modified by lensing, in the following way,

$$s^i_{\mathrm{obs}} = s^i_{\mathrm{int}} \left(1 + \alpha_s \kappa\right), \tag{4.53}$$

for some coefficient α_s which is unity in pure weak lensing theory. Denoting the number of galaxies for which sizes measurements are available by n_s, and the mean intrinsic size of the sample by $\bar{s}_{\mathrm{int}}$, the spectrum of the size field reduces, under the assumption of uncorrelated intrinsic sizes, to

$$C^s(l) = \bar{n}_s^2 \bar{s}^2_{\mathrm{int}} \alpha_s^2 C^\kappa(l) + \bar{n}_s \sigma_s^2. \tag{4.54}$$

The vector $\mathbf{v}$ and matrix N are now numbers, that are read out from the above equation, to be

$$\begin{aligned} v &= \bar{n}_s \bar{s}_{\mathrm{int}} \alpha_s, \\ N &= \bar{n}_s \sigma_s^2. \end{aligned} \tag{4.55}$$

leading to the effective noise

$$N^s_{\mathrm{eff}} = \frac{1}{\alpha_s^2} \left(\frac{\sigma_s}{\bar{s}_{\mathrm{int}}}\right)^2 \frac{1}{\bar{n}_s} \tag{4.56}$$

Second order, flexion Denoting with $\bar{n}_{\mathcal{F}}$ and $\bar{n}_{\mathcal{G}}$ the number of galaxies for which $\mathcal{F}$ and $\mathcal{G}$ are measured, the vectors linking the flexion to convergence are

$$\mathbf{v}^{\mathcal{F}} = -il\bar{n}_{\mathcal{F}} \begin{pmatrix} \cos\phi_l \\ \sin\phi_l \end{pmatrix} \tag{4.57}$$

and

$$\mathbf{v}^{\mathcal{G}} = -il\bar{n}_{\mathcal{G}} \begin{pmatrix} \cos 3\phi_l \\ \sin 3\phi_l \end{pmatrix}. \tag{4.58}$$

Using again the assumption of uncorrelated intrinsic components, we have the four dimensional diagonal noise matrix

$$N = \begin{pmatrix} \bar{n}_{\mathcal{F}}\sigma^2_{\mathcal{F}} \cdot 1_{2x2} & 0 \\ 0 & \bar{n}_{\mathcal{G}}\sigma^2_{\mathcal{G}} \cdot 1_{2x2} \end{pmatrix}, \tag{4.59}$$

leading to the effective noise, this time mode-dependent,

$$\frac{1}{N^{\mathcal{FG}}_{\text{eff}}} = l^2 \left(\frac{\bar{n}_F}{\sigma^2_F} + \frac{\bar{n}_G}{\sigma^2_G} \right). \tag{4.60}$$

5 On the use of Gaussian likelihoods for power spectra estimators

The text of this chapter follows very closely that of Carron (2012a).

In this note we revisit the Fisher information content of cosmological power spectra of Gaussian fields, when based on the assumption of a multivariate Gaussian likelihood for estimators, in order to comment on that assumption. We discuss that despite the fact that the assumption of a Gaussian likelihood is motivated by the central limit theorem, it leads if used consistently to a Fisher information content that violates the Cramér-Rao inequality, due to the presence of independent information from the parameter dependent covariance matrix. At any fixed multipole, this artificial term is shown to become dominant in the limit of a large number of correlated fields. While the distribution of the estimators does indeed tend to a Gaussian with a large number of modes, it is shown, however, that its Fisher information content does not, in the sense that their covariance matrix never carries independent information content. The reason why the information content of the spectra is correctly described by the usual formula (i.e. without the covariance term) in this estimator perspective is precisely the fact the the estimators have a chi-squared like distribution, and not a Gaussian distribution. The assumption of a Gaussian estimators likelihood is thus from the point of view of the information neither necessary nor really adequate, and we warn against the use of Gaussian likelihoods with parameter dependent covariance matrices for parameter inference from such spectra.

5.1 Introduction

Starting from the second half of the nineties (Jungman et al., 1996a,b; Tegmark, 1997; Tegmark et al., 1997), the calculation of Fisher information matrices in order to under-

stand quantitatively the constraining power of an experiment has become ubiquitous in cosmology, with its fundamental aspects now covered in cosmological textbooks (Dodelson, 2003; Durrer, 2008, section 11 and 6 respectively, e.g.), or (Heavens, 2009). This is especially true for experiments aimed at measuring power spectra of close to Gaussian fields, since in this case very handy analytical expressions can be obtained that can be applied in a variety of major cosmological subfields, such as for instance the CMB, galaxy clustering, weak lensing as well as their combination.
Nevertheless, even applied to Gaussian variables, Fisher information matrices are not totally exempt of subtleties. In this note, we revisit the two different possible perspectives on the Fisher information content of such spectra. One starting point is often the assumption of Gaussian errors. We point out that this assumption of a multivariate Gaussian likelihood for power spectra estimators is not fully consistent for the purpose of understanding their information content, due to a term violating the Cramér-Rao inequality, that we show is not necessarily small. Too much information is therefore assigned to the spectra under this assumption. We show that we can understand why this term is artificial precisely from the non Gaussian properties of the estimators, and discuss the reasons why the usual formula, i.e. without this term, or setting the covariance matrix to be parameter independent, still gives the correct amount of information.

The note is built as follows : In section 5.2 the two common approaches to the information content of spectra are discussed in details in the case of a single field. We clarify to what extent and why one is actually flawed, which is the source our comments on the use of Gaussian likelihoods for spectra. In section 5.3 we then turn to a correlated family of fields, where the violation of the Cramér-Rao inequality is shown to become substantial. We summarize and conclude in section 5.4.

We recall first the specific form of the Fisher information matrix, defined for a probability density function p as $F_{\alpha\beta} = \langle \partial_\alpha \ln p \, \partial_\beta \ln p \rangle$, α, β model parameters of interest, in the particular case of a multivariate Gaussian distribution with mean vector μ and covariance matrix Σ, (Tegmark et al., 1997; Vogeley and Szalay, 1996)

$$F_{\alpha\beta} = \sum_{i,j} \frac{\partial \mu_i}{\partial \alpha} \Sigma^{-1}_{ij} \frac{\partial \mu_j}{\partial \beta} + \frac{1}{2} \mathrm{Tr} \left[\Sigma^{-1} \frac{\partial \Sigma}{\partial \alpha} \Sigma^{-1} \frac{\partial \Sigma}{\partial \beta} \right] . \tag{5.1}$$

Remember that the Fisher information matrix has all the properties a meaningful measure of information on parameters must have, most importantly for us here the fact that any transformation of the data can only decrease its Fisher information matrix.

5.2 One field, gamma distribution

Consider a zero mean isotropic homogeneous Gaussian random field, in euclidean space or on the sphere. It is well known that the Gaussianity of the field is equivalent to the fact that the Fourier or spherical harmonic coefficients are independent complex Gaussian variables, only constrained by the reality condition. Equivalently, the real and imaginary

parts of those coefficients form independent real Gaussian variables. Such fields are entirely described by their spectrum, and so the extraction of the spectrum from the data with the help of an estimator is a fairly natural way to proceed for inference on parameters of interest. We place ourselves on the sphere, adopting the spherical harmonic notation for convenience. With the set of a_{lm} the harmonic coefficients, the model parameter dependent spectrum C_l is defined as

$$\langle a_{lm} a^*_{l'm'} \rangle = \delta_{ll'} \delta_{mm'} C_l. \tag{5.2}$$

Standard, unbiased quadratic estimators can be written as a sum over the number of Gaussian modes available, as

$$\hat{C}_l = \frac{1}{2l+1} \sum_{m=-l}^{l} |a_{lm}|^2. \tag{5.3}$$

We do not consider any source of observational noise, incomplete coverage or any other such issue, which are irrelevant for the points of our discussion. At this point, there are two ways to approach the problem of evaluating its information content in the cosmological literature. The first, - let us call this approach the 'field perspective' -, first calculates the information content of the field itself (equal to that of the set of a_{lm}'s), and then interprets this information as being the one of the spectrum. In this case, the information in the field is given by formula (5.1), with zero mean vector and diagonal covariance matrix C_l.

$$F_{\alpha\beta} = \frac{1}{2} \sum_{l=0}^{\infty} (2l+1) \frac{1}{C_l} \frac{\partial C_l}{\partial \alpha} \frac{1}{C_l} \frac{\partial C_l}{\partial \beta}, \tag{5.4}$$

where the factor $2l+1$ accounts for the number of independent Gaussian variables at a given multipole l. The sum is in practice restricted to the multipole range that will actually be measured to obtain the information in the spectrum to be extracted. A very small sample of works in this approach are Bernstein (2009); Hu and Jain (2004); Tegmark et al. (1997). This approach is arguably conceptually appealing, as it deals with the information content of the field itself, and does not require the definition of estimators nor the calculation of their covariance. However, for the same reasons, it is only indirectly connected to data analysis as it is not yet specified precisely how this information content is to be extracted. In the second approach - that we call the 'estimator perspective' - is defined first an estimator $\hat{C}_l$ for each C_l to be extracted, within some $\mathrm{l_{min}}$ and $\mathrm{l_{max}}$ (maybe with some bandwidth that we ignore here), and its covariance matrix $\Sigma_{ll'} = \left\langle \hat{C}_l \hat{C}_{l'} \right\rangle - \left\langle \hat{C}_l \right\rangle \left\langle \hat{C}_{l'} \right\rangle$ is calculated. Then it is argued that due to the central limit theorem, the distribution of the estimator will be approximately Gaussian. In the case of spectra of Gaussian fields, this is very well founded, at least for small scales modes, since (5.3) is a large sum of well behaved identically distributed independent variables. Then, under this assumption of Gaussianity, their information content is given by equation (5.1) with mean vector this time the set of C_l's itself and (model parameter dependent) covariance matrix $\Sigma_{ll'}$,

$$F_{\alpha\beta} = \sum_{l,l'=\mathrm{l_{min}}}^{\mathrm{l_{max}}} \frac{\partial C_l}{\partial \alpha} \Sigma^{-1}_{ll'} \frac{\partial C_{l'}}{\partial \beta} + \frac{1}{2} \mathrm{Tr} \left[\Sigma^{-1} \frac{\partial \Sigma}{\partial \alpha} \Sigma^{-1} \frac{\partial \Sigma}{\partial \beta} \right]. \tag{5.5}$$

It is well known that for the estimator (5.3) we have $\Sigma_{ll'} = \delta_{ll'} 2C_l^2/(2l+1)$. The Fisher information matrix, in the estimator perspective, reduces thus to

$$\begin{aligned} F_{\alpha\beta} = &\frac{1}{2}\sum_{l=\mathrm{l}_{\min}}^{\mathrm{l}_{\max}} (2l+1)\frac{1}{C_l}\frac{\partial C_l}{\partial\alpha}\frac{1}{C_l}\frac{\partial C_l}{\partial\beta} \\ &+ \frac{1}{2}\sum_{l=\mathrm{l}_{\min}}^{\mathrm{l}_{\max}} 4\frac{1}{C_l}\frac{\partial C_l}{\partial\alpha}\frac{1}{C_l}\frac{\partial C_l}{\partial\beta}. \end{aligned} \tag{5.6}$$

Clearly, the first term in the estimator perspective corresponds to that of the field perspective. However, the second term, coming from the derivative of the covariance matrix, is new. That term is not enhanced by a $(2l+1)$ factor, and is therefore very subdominant at high l. It is either usually neglected, or the covariance matrix of the estimators is inconsistently taken to be parameter independent, and in these cases the two approaches give the same results. Some expositions using explicitly this perspective include (Seo and Eisenstein, 2003, 2007; Tegmark, 1997), where the additional term is neglected, or the approach in (Dodelson, 2003, section 11.4.3), where the covariance matrix is treated as parameter independent. A work where this term plays a direct role is (Eifler et al., 2009), where the authors specifically study the impact of parameter dependent covariance matrices for parameter estimation using such Gaussian likelihoods.

Beyond the question of the quantitative relevance of this additional term, its very appearance is however very disturbing. Under this arguably reasonable Gaussian assumption, our estimator (5.3) is found to carry more information than the full field, even on the smallest scales. This obviously violates the most fundamental property of Fisher information, i.e. that information can only be at best conserved when transforming the data (in this case reducing the field to its spectrum), a fact essentially equivalent to the celebrated Cramér-Rao inequality (Tegmark et al., 1997). Something must clearly have gone wrong in the assumption of a Gaussian likelihood for our spectra.

To understand what has happened, it is worth tracking the exact distribution and information content of the estimator (5.3). Since they are independent at different l, we can work at a fixed l, and the total information content of these estimators will simply be the sum over l of the information of the estimator at fixed l. Under our assumptions, the estimator is a sum of squares of $2l+1$ independent Gaussian variables, and its probability density function can be obtained with no difficulty. The exact distribution is the gamma probability density function with shape parameter k and location parameter θ as follows

$$p(\hat{C}_l|\alpha,\beta) = \exp\left(-\hat{C}_l/\theta\right)\frac{\hat{C}_l^{k-1}}{\theta^k\Gamma(k)}, \tag{5.7}$$

with

$$k := \frac{1}{2}(2l+1), \quad \theta(\alpha,\beta) := \frac{2C_l}{2l+1}, \tag{5.8}$$

and where Γ is the gamma function. It is well known that the gamma distribution does indeed tend towards the Gaussian distribution for large k, with mean $\mu = k\theta = C_l$ and variance $\sigma^2 = k\theta^2 = 2C_l^2/(2l+1)$, as expected. However, its Fisher information content

does not tend to that of the Gaussian. In our case, since only θ is parameter dependent, we have that the Fisher information in the estimator density function (5.7) is

$$F^l_{\alpha\beta} = \frac{\partial\theta}{\partial\alpha}\frac{\partial\theta}{\partial\beta}\left\langle\left(\frac{\partial \ln p(\hat{C}_l)}{\partial\theta}\right)^2\right\rangle. \tag{5.9}$$

Since $\partial_\theta \ln p = (\hat{C}_l - k\theta)/\theta^2$, and $\partial_\alpha\theta = 2\theta\partial_\alpha C_l/C_l$, we obtain with straightforward algebra

$$F^l_{\alpha\beta} = \frac{1}{2}(2l+1)\frac{1}{C_l}\frac{\partial C_l}{\partial\alpha}\frac{1}{C_l}\frac{\partial C_l}{\partial\beta}. \tag{5.10}$$

Summing over l, we recover the first term of (5.6), but not the second. We have recovered the field perspective result (5.4) at any l without the Gaussian assumption but with the exact distribution. It turns out that even though the variance of the gamma distribution is parameter dependent, it does not in fact contribute to the information. This can be seen as the following. Consider the information in the mean only of the estimator. From the Cramér-Rao inequality this must be less than the total information,

$$\frac{1}{\sigma^2}\frac{\partial\mu}{\partial\alpha}\frac{\partial\mu}{\partial\beta} \leq F^l_{\alpha\beta}. \tag{5.11}$$

Plugging in the values for the mean and variance leads in fact to the result that the inequality is an equality, so that the mean of the estimator captures all of its information.

In summary, the Gaussian approximation assumes the mean and the variance of the estimator are uncorrelated, such that both contributes to the information, while for the exact gamma, they are degenerate in such a way that the variance does not carry independent information. Another way to see this, that we will use below when the exact form of the distribution will be less convenient, comes from the fact that $\partial_\theta \ln p(\hat{C}_l)$ is a first order polynomial in $\hat{C}_l$. It can be shown namely that the first n moments capture all the information precisely when $\partial_\alpha \ln p$ is a polynomial of order n (Carron, 2011).

5.3 Several fields, Wishart distribution

It is instructive to see how these considerations generalize to a situation of a family of n jointly zero mean Gaussian correlated fields, where the analysis proceeds through the extraction of the spectra and cross spectra. In this case, the C_l of the above discussion becomes a $n \times n$ (possibly complex) Hermitian matrix

$$\left\langle a^i_{lm} a^{j*}_{lm}\right\rangle = \delta_{ll'}\delta_{mm'}C^{ij}_l, \quad C^\dagger = C. \tag{5.12}$$

From the hermiticity property there are only $n(n+1)/2$ non redundant spectra. Adequate estimators are defined by a straightforward generalization of equation (5.3),

$$\hat{C}^{ij}_l = \frac{1}{2l+1}\sum_{m=-l}^{l} a^i_{lm}a^{j*}_{lm}. \tag{5.13}$$

While the estimators are still independent for different l's, the different components at a given l are not. The information content of the set of a^i_{lm} in the field perspective is still given by formula (5.1) for zero mean Gaussian variables. Explicitly, at a given l,

$$F^l_{\alpha\beta} = \frac{1}{2}(2l+1)\mathrm{Tr}\left[C_l^{-1}\frac{\partial C_l}{\partial\alpha}C_l^{-1}\frac{\partial C_l}{\partial\beta}\right]. \tag{5.14}$$

In the estimator perspective, assuming the estimators $\hat{C}^{ij}_l, i \leq j$ are jointly Gaussian, we have instead

$$\sum_{i<j,k<l=1}^{n} \frac{\partial C^{ij}_l}{\partial\alpha}\Sigma^{-1}_{ij,kl}\frac{\partial C^{kl}_l}{\partial\beta} + \frac{1}{2}\mathrm{Tr}\left[\Sigma^{-1}\frac{\partial\Sigma}{\partial\alpha}\Sigma^{-1}\frac{\partial\Sigma}{\partial\beta}\right], \tag{5.15}$$

where the covariance matrix is

$$\begin{aligned}\Sigma_{ij,kl} &= \left\langle \hat{C}^{ij}_l \hat{C}^{kl}_l \right\rangle - C^{ij}_l C^{kl}_l \\ &= \frac{1}{2l+1}\left(C^{ik}_l C^{jl}_l + C^{il}_l C^{jk}_l\right).\end{aligned} \tag{5.16}$$

While it may not be immediately obvious this time, it has been noted (Hu and Jain, 2004, e.g.) that the first term in (5.15) is rigorously equivalent to the expression from the field perspective (5.14). The estimator perspective under the assumption of a multivariate Gaussian distribution for $\hat{C}_l$ thus still violates the Cramér-Rao inequality due the presence of the second term. Since since this term is not enhanced by a factor of $2l+1$ we expect it to be subdominant again. However, it is less true this time than in the one dimensional setting : using the explicit form of the inverse covariance matrix,

$$\begin{aligned}\Sigma^{-1}_{ij,kl} = (2l+1)&\left(C_l^{-1,ik}C_l^{-1,jl} + C_l^{-1,il}C_l^{-1,jk}\right) \\ &\cdot\left(1 - \frac{1}{2}\left(\delta_{ij}+\delta_{kl}\right) + \frac{1}{4}\delta_{ij}\delta_{kl}\right),\end{aligned} \tag{5.17}$$

one can derive with some lengthy but straightforward algebra the following expression for the violating term,

$$\begin{aligned}&\frac{1}{2}\mathrm{Tr}\left[\Sigma^{-1}\frac{\partial\Sigma}{\partial\alpha}\Sigma^{-1}\frac{\partial\Sigma}{\partial\beta}\right] \\ &= \frac{1}{2}(n+2)\mathrm{Tr}\left[C_l^{-1}\frac{\partial C_l}{\partial\alpha}C_l^{-1}\frac{\partial C_l}{\partial\beta}\right] \\ &\quad + \frac{1}{2}\mathrm{Tr}\left[C_l^{-1}\frac{\partial C_l}{\partial\alpha}\right]\mathrm{Tr}\left[C_l^{-1}\frac{\partial C_l}{\partial\beta}\right],\end{aligned} \tag{5.18}$$

for any number n of fields. If $n = 1$, we recover indeed (5.6). While the term is still subdominant at high l, the situation is yet a bit less comfortable. The number of fields is not necessary very small in cosmologically relevant situations, such as tomographic joint shear and galaxy densities analysis in redshift slices, to which one may also add magnification, flexion fields, etc. Writing schematically $n = N_f N_{\mathrm{bin}}$, where N_{bin} is the number of bins and N_f the number of fields per bin, we have e.g. $N_f = 3$ for the galaxy density and the two shear fields, $N_f = 4$ including magnification, $N_f = 8$ adding hypothetically the four

flexion fields, and so on. Comparing (5.14) and (5.18), and neglecting the second term in (5.18), we have that at

$$l \sim \frac{1}{2} N_f N_{\rm bin}, \tag{5.19}$$

the Cramér-Rao violating term is actually still the dominant one. Note that this is still optimistic. Due to the product of two traces in the second term in (5.18), one can expect roughly the same scaling with n as the first term. Thus, the correct l in (5.19) may generically be closer to

$$l \sim N_f N_{\rm bin}. \tag{5.20}$$

From the discussion in section 5.2, we can easily guess what went wrong. Consider the information content of the means of the estimators exclusively. This is given for any probability density function by weighting the derivatives of the means with the inverse covariance matrix, and is thus equal to the correct, first term in (5.15). Since already the means of the estimators do exhaust the information in the field, we can therefore already conclude that the total information content of the estimators must be equal to that of their means, and in particular that the covariance does not contribute to the information. As before, the second term in the estimator perspective is an artifact of the Gaussian assumption. It is interesting though to derive as above more explicitly why only the means carry information, from the shape of the joint probability density of the estimators. The remainder of this section sketches how this can be simply performed, leading to equation (5.26).

We restrict ourselves now for the sake of notation to the case of two fields, $n = 2$, but the following argumentation holds for any n. The exact joint distribution for the three estimators $\hat{C}_l = (\hat{C}_l^{11}, \hat{C}_l^{12}, \hat{C}_l^{22})$, is given from the rules of probability theory as

$$p(\hat{C}_l|\alpha,\beta) = \left\langle \prod_{i\leq j=1}^{2} \delta^D \left(\hat{C}_l^{ij} - \frac{1}{2l+1} \sum_{m=-l}^{l} a_{lm}^i a_{lm}^{j*} \right) \right\rangle \tag{5.21}$$

where δ^D is the Dirac delta function. The average is over the joint probability density for the two sets of harmonic coefficients a_{lm}^1 and a_{lm}^2. Define the vector

$$\mathbf{a}_l = (a_{l-l}^1, \cdots, a_{ll}^1, a_{l-l}^2, \cdots, a_{ll}^2). \tag{5.22}$$

Since the a_{lm} are zero mean Gaussian variables with correlations as given in (5.12), this probability density function is given by

$$\frac{1}{Z(\alpha,\beta)} \exp\left(-\frac{1}{2} \mathbf{a}_l^\dagger \cdot \mathbf{C}_l^{-1} \mathbf{a}_l \right), \tag{5.23}$$

with

$$\mathbf{C}_l = \begin{pmatrix} C_l^{11} \cdot 1_{2l+1} & C_l^{12} \cdot 1_{2l+1} \\ C_l^{21} \cdot 1_{2l+1} & C_l^{11} \cdot 1_{2l+1} \end{pmatrix}, \tag{5.24}$$

where 1_{2l+1} is the unit matrix of size $2l+1$. $Z(\alpha,\beta)$ is the normalization of the density for $\mathbf{a}$, that does depend on the model parameters through the determinant of the $\mathbf{C_l}$ matrix.

The inverse matrix $\mathbf{C_l^{-1}}$ has the same block structure, with entries being those of C_l^{-1}. In the following we are not really interested in keeping track of the exact value of the components of this matrix, but only that they are dependent on the model parameters. With the understanding that $C_l^{-1} =: D_l$, we have thus, due to the sparse structure of the $\mathbf{C}_l^{-1}$ matrix and the Dirac delta functions in (5.21),

$$-\frac{1}{2}\mathbf{a}_l^\dagger \cdot \mathbf{C}_l^{-1}\mathbf{a}_l = -\frac{1}{2}(2l+1)\sum_{i,j=1,2} D_l^{ij}\hat{C}_l^{ij} \tag{5.25}$$

Due to the presence of the Dirac delta functions, we can thus take the exponential (5.23) out of the integral in (5.21). Writing explicitly the dependency of the different terms on $\hat{C}_l$ and the model parameters, we obtain the following form

$$\begin{aligned} &p(\hat{C}_l|\alpha,\beta) \\ &= \frac{f(\hat{C}_l)}{Z(\alpha,\beta)}\exp\left(-\frac{1}{2}(2l+1)\sum_{i,j=1,2} D_l^{ij}(\alpha,\beta)\hat{C}_l^{ij}\right) \end{aligned} \tag{5.26}$$

which generalizes the gamma distribution, equation (5.7), in this multidimensional case. The factor $f(\hat{C}_l)$ is what is left from the integral (5.21) when the density for the set of a_{lm} is taken out, i.e. the volume of the space spanned by the a_{lm}'s that satisfies the constraints set by the Dirac delta function. It is thus a factor that depends on $\hat{C}_l$ but importantly for us not on the model parameters [1]. The point of the representation (5.26) is that it is immediate that $\partial_\alpha \ln p(\hat{C}_l)$ is a polynomial first order in the components of $\hat{C}_l$. Second order terms, corresponding to information within the covariance matrix never appear, however close to a Gaussian the exact density function might be. It follows that the total Fisher information matrix is always equal to that of the mean, even if we did not derive the exact shape of the distribution.

5.4 Summary and conclusions

We discussed two common perspectives (the 'field' and 'estimator' perspectives) on the Fisher information content of cosmological power spectra, and why in the estimator perspective the assumption of a Gaussian likelihood of the spectra estimators violates the Cramér-Rao inequality, assigning the estimators more information than there is in the full underlying fields. Under the assumption of Gaussianity of the estimators, their means and covariance matrix are artificially rendered uncorrelated, creating an additional piece of information in their covariance, that we showed was inexistent by calculating the exact information content of the estimators true probability density function. We showed that this violating term can become dominant in the limit of a large number of fields. Using Gaussian likelihoods consistently, i.e. with parameter dependent covariance matrices, as argued for example in (Eifler et al., 2009), assigns therefore far too much information to

[1] The prefactors in (5.26) can be obtained in closed form, leading to the Wishart density function. See (Hamimeche and Lewis, 2008, e.g.)

the spectra in this regime, and should thus be avoided.

In the estimator perspective of the derivation of the Fisher information matrix, this term is usually neglected. This note clarifies why it should not be present in the very first place, and how the agreement between the field and estimator perspective can thus arguably be seen as an happy cancellation of two inconsistencies. It is interesting to note that the reason why we still find the exact result in the estimator perspective without this wrong piece is that this expression is also the exact Fisher information content of the exact, for low l strongly non Gaussian, distribution of the estimators, the central limit theorem playing actually no role.

The other lesson we can take from this work is that in general, when in doubt about the joint distribution of a set of estimators, a safe choice of information content is always that of their means exclusively, which requires only the knowledge of their covariance. Provided the covariance matrix is correctly chosen, one is indeed sure for any probability density function from the properties of Fisher information to make a conservative evaluation, that does not rely on any further assumptions on its shape. Thus, leaving apart the question of the very accuracy of the approximation itself, using a Gaussian likelihood with parameter independent covariance matrix, having the entire information in the means, while not entirely consistent remains a safe prescription in the sense that a conservative information content is always assigned to the estimators.

6 N-point functions in lognormal density fields

In this chapter we discuss extensively the information content of the lognormal field as a model for the matter fluctuation field in cosmology. It is built out of published as well as yet unpublished elements. The text in sections 6.1 and 6.4 follows closely that of Carron (2011), and that of sections 6.1.1, 6.1.2, 6.3 and 6.5 that of Carron and Neyrinck (2012). On the other hand, sections 6.2, 6.6 and 6.7 present unpublished material.

6.1 Introduction

The cosmological matter density field is becoming more and more directly accessible to observations with the help of weak lensing (Bartelmann and Schneider, 2001; Munshi et al., 2006; Refregier, 2003; Schneider et al., 1992). Its statistical properties are the key element in trying to optimize future large galaxy surveys aimed at answering actual fundamental cosmological questions, such as the nature of the dark components of the universe (Caldwell and Kamionkowski, 2009; Frieman et al., 2008). To this aim, Fisher's measure of information on parameters (Fisher, 1925; Rao, 1973; van den Bos, 2007) has naturally become of standard use in cosmology. It provides indeed an handy framework, in which it is possible to evaluate in a quantitative manner the statistical power of some experiment configuration aimed at some observable (Albrecht et al., 2006; Amara and Réfrégier, 2007; Bernstein, 2009; Hu and Jain, 2004; Hu and Tegmark, 1999; Parkinson et al., 2007; Tegmark, 1997; Tegmark et al., 1997, e.g.). Such studies are in the vast majority of cases limited to Gaussian probability density functions, or perturbations therefrom, and deal mostly with the prominent members of the correlation function hierarchy (Peebles, 1980), or equivalently their Fourier transforms the polyspectra, such as the matter power spectrum.

The approach via the correlation function hierarchy is very sensible in the nearly linear regime for at least two reasons. First, in principle, the correlations are the very elements that cosmological perturbation theory is able to predict in a systematic manner (see Bernardeau et al. (2002) for a review, or the more recent (Matsubara, 2011) and the numerous references in it). Second, primordial cosmological fluctuations fields are believed to be accurately described by the use of Gaussian statistics. It is well known that the correlations at the two-point level provide a complete description of Gaussian fields. It is therefore natural to expect this approach to be adequate throughout the linear and the mildly non linear regime, when departures from Gaussianity are small.

Deeper in the non linear regime, fluctuations grow substantially in size, and tails in the matter probability density function do form. A standard prescription for the statistics of the matter field in these conditions is the lognormal distribution, various properties of which are discussed in details in an astrophysical context in (Coles and Jones, 1991). It was later shown to be reproduced accurately, both from the observational point of view as well as in comparison to standard perturbation theory and N-body simulations (Bernardeau, 1994; Bernardeau and Kofman, 1995; Kayo et al., 2001; Taylor and Watts, 2000; Wild et al., 2005), in low dimensional settings. The lognormal assumption is also very much compatible with numerical works (Neyrinck et al., 2009, 2011) showing that the spectrum of logarithm of the field $\ln 1 + \delta$ carries much more information than the spectrum of δ itself. The first evaluation of the former within the framework of perturbation theory appeared recently (Wang et al., 2011).

Lognormal statistics (Aitchison and Brown, 1957, for a textbook presentation) are not innocuous. More specifically, the lognormal distribution is only one among many distributions that leads to the very same series of moments. This fact indicates that, going from the distribution to the moments, one may be losing information in some way or another. A fundamental limitation of the correlation function hierarchy in extracting the information content of the field in the non linear regime could therefore exist, if its statistics are indeed similar to the lognormal. This important fact was already mentioned qualitatively in (Coles and Jones, 1991), but it seems no quantitative analysis is available at present.

In this chapter we provide first answers to these issues, looking at the details of the structure of the information within the lognormal field and its N-point moments. This chapter is built as follows :

We start by exposing briefly two fundamental limitations of the lognormal assumption in section 6.2. We continue in section 6.3 defining explicit families of fields that have the very same hierarchy of N-point moments than the lognormal field. This is making obvious that the hierarchy does not provide a complete description of the field, and therefore does not contain the entire information. In section 6.4, we quantify the importance of this aspect in terms of Fisher information : we solve in that section for the statistical power of the moments of the lognormal distribution exactly at all orders, and for their efficiency in capturing information. We also compare these predictions to standard perturbation theory. In section 6.5 we then use these results derived in section 6.4 to make

a successful connection to the N-body simulation results mentioned above, that showed that the statistical power of the spectrum of the logarithm of the field is larger than that of the original fluctuation field. In section 6.6, we extend the results of 6.4 by allowing parameters to create correlations between the variables, and discuss in this light the fact that the improvement seen in simulations are seen to be mostly parameter independent. We conclude in section 6.7 with a remark on the derivation of the statistical power of the N-point moments for an arbitrary lognormal field.

6.1.1 Notation and conventions

We will be dealing throughout this chapter with random vectors $\rho = (\rho_1, \cdots, \rho_d)$, being the sample of a density field ρ,

$$\rho_i = \rho(x_i) > 0. \tag{6.1}$$

We place ourselves in 3-dimensional cartesian space for convenience. For a vector $\mathbf{n} = (n_1, \cdots, n_d)$ of non negative integers (multiindex), we write as in chapter 3 $\rho^{\mathbf{n}}$ the monomial in d variables,

$$\rho^{\mathbf{n}} = \rho(x_1)^{n_1} \cdots \rho(x_d)^{n_d}. \tag{6.2}$$

Throughout this chapter, we reserve bold letters for vectors of integers exclusively.

Let $p_\rho(\rho)$ be a d-dimensional probability density function such that all correlations of the form $\langle \rho^{\mathbf{n}} \rangle$ are finite. We write the moment $\langle \rho^{\mathbf{n}} \rangle$ with $m_{\mathbf{n}}$. Explicitly,

$$m_{\mathbf{n}} = \left\langle \rho^{n_1}(x_1) \cdots \rho^{n_d}(x_d) \right\rangle. \tag{6.3}$$

Correlations of order n are given by moments such that the order $|\mathbf{n}|$ of the multiindex, defined as

$$|\mathbf{n}| := \sum_{i=1}^{d} n_i \tag{6.4}$$

is equal to n. These moments coincide with the values of a continuous n-point correlation function on the grid sampled by $(x_1, \cdots, x_d)$. We write δ for the dimensionless fluctuation field, and A for the field defined by $\ln \rho$:

$$A := \ln \rho, \quad \delta := \frac{\rho - \bar{\rho}}{\bar{\rho}}. \tag{6.5}$$

Such assignments involving ratios or logarithms of d-dimensional quantities should always be understood component per component.

6.1.2 Definition and basic properties of correlated lognormal variables

We say the d-dimensional vector $\rho := (\rho(x_1), \cdots, \rho(x_d))$ is lognormal if the d-dimensional probability density function for A is Gaussian. Explicitly,

$$p_A(A) = \frac{1}{(2\pi|\xi_A|)^{d/2}} \exp\left(-\frac{1}{2}(A - \bar{A}) \cdot \xi_A^{-1}(A - \bar{A}) \right), \tag{6.6}$$

where $\bar{A}$ is the mean vector of A, and ξ_A its covariance matrix,

$$[\xi_A]_{ij} = \left\langle \left(A(x_i) - \bar{A}(x_i)\right) \left(A(x_j) - \bar{A}(x_j)\right) \right\rangle . \tag{6.7}$$

The probability density for the vector ρ itself is then by construction a d-dimensional lognormal distribution. We name it for further reference as p_ρ^{LN}. From the rules of probability theory holds

$$p_\rho^{LN}(\rho) = \frac{p_A(\ln \rho)}{\prod_{i=1}^d \rho(x_i)}. \tag{6.8}$$

The means and two point correlations of A and δ are in one to one correspondence. We have

$$\bar{A} = \ln \bar{\rho} - \frac{1}{2}\sigma_A^2 \tag{6.9}$$

where σ_A^2 is the diagonal of ξ_A, i.e. the variances of the individual d points. Also,

$$[\xi_A]_{ij} = \ln\left(1 + [\xi_\delta]_{ij}\right), \quad [\xi_\delta]_{ij} := \langle \delta(x_i)\delta(x_j)\rangle . \tag{6.10}$$

Especially, the variances are related through

$$\sigma_A^2 = \ln\left(1 + \sigma_\delta^2\right). \tag{6.11}$$

In the multiidindex notation, the N-point moments of ρ take the following simple form,

$$m_{\mathbf{n}} = \exp\left(\bar{A}\cdot\mathbf{n} + \frac{1}{2}\mathbf{n}\cdot\xi_A\mathbf{n}\right). \tag{6.12}$$

The mapping from A to ρ is invertible, and A is a multivariate Gaussian distribution. Therefore the total Fisher information content of ρ is given by

$$F_{\alpha\beta} = \frac{\partial \bar{A}^T}{\partial\alpha}\xi_A^{-1}\frac{\partial\bar{A}}{\partial\beta} + \frac{1}{2}\mathrm{Tr}\left[\frac{\partial\xi_A}{\partial\alpha}\xi_A^{-1}\frac{\partial\xi_A}{\partial\beta}\xi_A^{-1}\right]. \tag{6.13}$$

6.2 Fundamental limitations of the lognormal assumption

To open this chapter, we discuss two fundamental limitations of the lognormal assumption for the matter density field. We are interested in the validity of the approximation, i.e. the mere possibility to describe at all the ΛCDM matter density field with an homogeneous isotropic lognormal field. It is indeed not guaranteed a priori that quantities derived from this assumption will be well defined. For instance, any sensible statistical model requires covariance and information matrices to be positive, and it is not clear yet whether the lognormal assumption on the matter density field satisfies these conditions.

It is important to note that this is very different from testing the accuracy of the approximation. For instance, given a two-point correlation function (or equivalently a positive

Fourier transform, the power spectrum), it is always possible to use the Gaussian field as a prescription, even though this approximation may be extremely inaccurate. This is because the two-point function is the only relevant ingredient, and the Gaussian field with that two-point correlation function is well defined statistically speaking in all cases.

With the lognormal field this is not the case anymore. Namely, given the two-point correlation function ξ_δ with positive Fourier transform, we see from (6.10), in a continuous notation, that there is the additional constraint that

$$\xi_A(r) = \ln\left(1 + \xi_\delta(r)\right) \tag{6.14}$$

must be a valid two-point correlation function as well, since it is that of the A field. Since A is by definition a Gaussian field, we see that formally for the lognormal assumption to be valid, the following Fourier transform must be positive for all k,

$$\int d^3r \ln(1 + \xi_\delta(r))e^{-ik\cdot r} = P_A(k) \overset{!}{\geq} 0. \tag{6.15}$$

We discuss in the following that these non trivial constraints are not satisfied both on the largest and on the smallest scales in our current understanding of the ΛCDM power spectrum.

6.2.1 Largest scales

An immediate issue is the power of the A field on the very largest scales. Setting $k = 0$ in (6.15) gives us

$$P_A(0) = \int d^3r \ \ln\left(1 + \xi_\delta(r)\right). \tag{6.16}$$

On the other hand, we have that for any argument x holds $\ln(1 + x) \leq x$, with equality if and only if $x = 0$. For this reason, we have

$$P_A(0) < \int d^3r \ \xi_\delta(r) = P_\delta(0). \tag{6.17}$$

Now, the spectrum of the fluctuation field is believed to have the scale free shape of a power law $P_\delta \propto k^n$, with n very close to unity (Komatsu et al., 2011). Extrapolated to the smallest wavenumbers, this functional form obviously assigns zero power at zero. We find therefore the contradiction

$$P_A(0) < 0. \tag{6.18}$$

On the largest scales, the power of the A field is thus found to be negative under this lognormal assumption, rendering it formally untenable.

Nevertheless, this issue cannot be considered a real shortcoming of the model. Namely, these large scales are in no way observable and are irrelevant to any realistic situation. In particular the power at zero, formally the variance of the mean of the field from realisations to realisations of the Universe within the statistical model carries no meaning for observations.

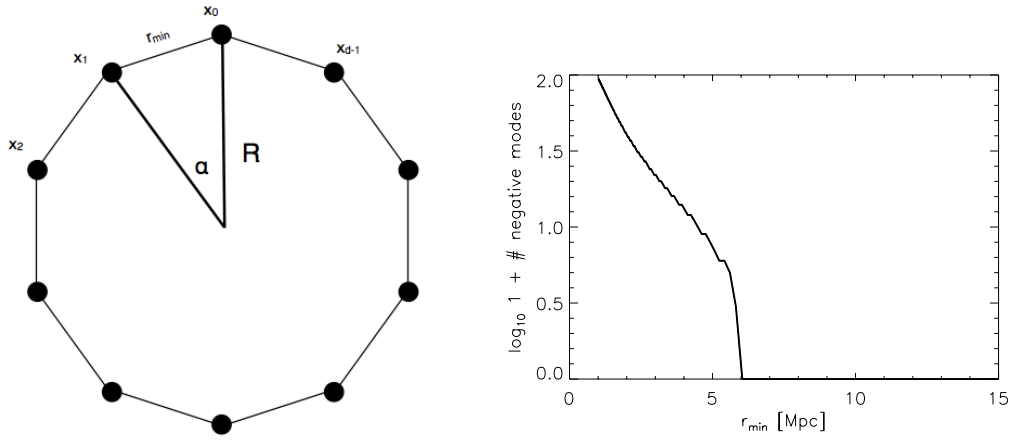

Fig. 6.1 — Left panel : The maximally symmetric configuration of points, including a finite range of distances, chosen to test the positivity of the spectrum of the A field, for the ΛCDM two-point function. Right panel : On a logarithmic scale, the number of wavenumbers such that the spectrum of the A field is found to be negative. Such mode appear once scales below $5-6$ Mpc are included, indicating that a lognormal description is no longer well defined below these scales.

6.2.2 Smallest scales

More interesting are the small scales, since these are the ones that are accessible to observations. Performing the angular integration in (6.15), the conditions that the two-point function must obey become

$$\int_0^\infty dr\, r^2 \ln\left(1+\xi_\delta(r)\right) j_0(kr) \geq 0, \text{ for all } k, \tag{6.19}$$

where $j_0(x) = \sin(x)/x$ is the first spherical Bessel function. These relations are neither convenient to test nor the most meaningful however. On one hand, due to the r^2 factor it requires the correlation function at very large argument. On the other hand, the inequalities (6.19) represent the constraints puts on a continuous lognormal field in an infinite volume. They are thus rather formal and do not correspond to a situation that can occur in practice. We therefore build another set of similar identities that are both simpler to test and more importantly where we can investigate a range of specific scales without ambiguity.

For this we use a finite but maximally symmetric configuration of points x_i where the corresponding constraints become the positivity of the two-point correlation matrix

$$[\xi_A]_{ij} = \left\langle \left(A(x_i) - \bar{A}\right)\left(A(x_j) - \bar{A}\right)\right\rangle \tag{6.20}$$

The ideal configuration for this purpose is the classic chain with d points described in figure 6.1.

In this configuration, periodic boundary conditions are built in by construction, such that we are still able to write the two-point correlation matrix as the transform of a positive spectrum. The two parameters R and the number d of points (or equivalently $r_{\min}$) allow us to keep complete control over the range of scales that are being tested. From rotational invariance, it is seen that the distances involved are $r_j, j = 0, \cdots d-1$, given by

$$r_j = |x_j - x_0| = 2R \sin\left(j\frac{\pi}{2d}\right), \quad j = 0, \cdots, d-1. \tag{6.21}$$

The correlation matrix becomes

$$[\xi_A]_{ij} = \xi_A(r_{|j-i|}). \tag{6.22}$$

This matrix has the structure of a so-called circulant matrix. Circulant matrices are always diagonalised by the discrete Fourier transform : the eigenvalues are

$$P_k = \sum_{j=0}^{d-1} \xi_A(r_j) \exp\left(-2\pi i k \frac{j}{N}\right), \quad k = 0, \cdots, d-1, \tag{6.23}$$

and the matrix can be written as

$$\xi_A(r_j) = \frac{1}{d} \sum_{k=0}^{d-1} P_k \exp\left(2\pi i k \frac{j}{N}\right). \tag{6.24}$$

The eigenvalues P_k must be positive for ξ_A to be a valid two-point correlation function.

We have thus found a new set of constraints :

$$\begin{gathered} 0 \leq \sum_{j=0}^{d-1} \ln\left(1 + \xi_\delta(r_j)\right) \exp\left(-2\pi i k \frac{j}{d}\right), \\ k = 0, \cdots, d-1, \text{ and for any } R \text{ and } d. \end{gathered} \tag{6.25}$$

These constraints can be trivially tested using an FFT algorithm, and involve distances between $r_{\min} \sim 2\pi R/d$ and $2R$ exclusively. A two-point correlation function ξ_δ must satisfy these constraints in order for the field to be possibly lognormal.

In the limit of large number of points, the sum goes over to an integral, and it is easy to show that these constraints become

$$0 \leq \frac{1}{\pi} \int_0^\pi d\alpha \;\; \ln\left[1 + \xi_\delta\left(2R \sin\left(\frac{\alpha}{2}\right)\right)\right] \cos\left(\alpha k\right), \tag{6.26}$$

again for $k = 0, 1, \cdots$ and any R. For very large R , P_k becomes the spectrum of the one-dimensional field.

We tested these inequalities for the ΛCDM correlation function. We generated the power spectrum with the help of the iCosmo package (Refregier et al., 2008) smoothed it with a top-hat filter of size $r_{\min}$, Fourier transformed it to obtain ξ_δ, and then tested straightforwardly equations (6.25) with an FFT algorithm.

While we found to the radius R to be of very little relevance, these inequalities break down on scales $r_{\rm min}$ below $5-6$ Mpc (corresponding to a variance σ_δ^2 of roughly 3) as shown on the right panel of figure 6.1, where the spectrum on these scales gets negative. It follows that the lognormal assumption is statistically speaking not well defined on these small scales for a ΛCDM cosmology.

6.3 Explicit fields with the same hierarchy of N-point functions

In this section, we present explicit families of density fields with the same N-point moments of any order that the lognormal field. This allows us to demonstrate in clearest possible way the fact that the hierarchy of N-point functions never provide a complete description of a lognormal field, and discuss in the light of these families some implications. We first begin in 6.3.1 with a discussion on the multidimensional moment problem and its link to some concepts already present in the cosmological literature. The families are then defined in 6.3.2 and 6.3.3. The appendix 6.3.4 contains the proofs of the key statements.

6.3.1 The problem with tailed fields

As already mentioned in chapter 3, in one dimension the fact that the hierarchy does not always specify fully the distribution is a well known and still active topic of research in the theory of moments in mathematics (Akhiezer, 1965; Shohat and Tamarkin, 1963; Simon, 1997, for classical references). The moment problem is to find a distribution corresponding to a given moment series. When a unique solution exists, it is called a determinate moment problem. When several exist (in this case always infinitely many), it is called an indeterminate moment problem. We can refer at this point to Coles and Jones (1991) for a discussion in a cosmological context. The theory of the moment problem in several dimensions is less developed, but typical criteria that guarantee determinacy, or indeterminacy, linked to the decay rate of the distribution, stay basically unchanged. Guiding us throughout the discussion in this section will be the following instance: for any dimension d, if

$$\left\langle e^{c|\rho|} \right\rangle < \infty, \quad |\rho| = \left(\rho_1^2 + \cdots + \rho_d^2\right)^{1/2} \tag{6.27}$$

for some $c > 0$, then the moment problem corresponding to the moments of that distribution is determinate (Dunkl and Xu, 2001, theorem 3.1.17). By a 'tailed' distribution, we have in mind in this work a decay at infinity which is less than exponential, and thus for which this criterion fails. In this regime, there may thus be several distributions with the same hierarchy of correlations.

It should be clear why this can have in general a dramatic impact for parameter inference from correlations. Imagine a series of distributions with identical correlations at all orders,

one of these distributions being the one that actually describes the observations. Since the distributions are different, they will make in general different predictions for observables other than the correlations. Pick for definiteness an observable $\langle f(\rho)\rangle(\alpha)$ with different predictions among this family of distribution, α any model parameter. The knowledge of the entire hierarchy is unable to distinguish from these different predictions for $\langle f\rangle$, since they result from equally valid distributions. If α enters the true distribution in such a way that it makes a sharp prediction on the value of $\langle f\rangle$, this is highly valuable information definitely lost to an analyst extracting correlations exclusively. On the other hand this argument allows us also to see that this effect can become relevant only when perturbation theory breaks down. If the fluctuation field δ is small, f can be expanded in powers of δ, and thus $\langle f\rangle$ can be obtained in an unique way from the correlation hierarchy of δ.

The formalism developed in chapter 3 provides us with a remarkable way to understand what is happening there in terms of Fisher information, familiar to cosmologists. In particular, we have seen that the distributions for which the Fisher information matrix is within the entire hierarchy are precisely those for which the functions $\partial_\alpha \ln p$ can be written as a power series over the range of p. If not, the mean squared residual to the best series expansion is the amount of Fisher information absent from the hierarchy. It is simple to show that criterion (6.27), that guarantees that the distribution is uniquely set by its correlations, implies as well that the entire amount of Fisher information is within the hierarchy : this follows from the very next theorem of the same reference (Dunkl and Xu, 2001, theorem 3.1.18), that states that the polynomials in the d variables form a dense set of functions with respect to the least mean squared residual criterion, if (6.27) is met. In particular the functions $\partial_\alpha \ln p$ can be arbitrarily well approximated by polynomials with respect to that criterion, and therefore the correlations contain all of the Fisher information.

It is important to note that if criterion (6.27) happens to be met due to a cutoff at a large value $\rho_{\rm cut}$, on a otherwise tailed distribution, the correlations still are poor probes for any practical purposes. For instance, if a variable is lognormal over a very long range, but decay quickly at infinity starting from $\rho_{\rm cut}$. Indeed, if $\rho_{\rm cut}$ is large enough, the correlations of order up to, say, $2N$, will be identical to that of the lognormal. Since the information content of the first N correlations depends on the first $2N$ only, they will be equally poor probes as for the lognormal (this will be quantified in the next two sections). They will contain the exact same amount of Fisher information as the ones of the lognormal. It is the correlations of order $> N$, that are able to feel the cutoff, that will make up for the difference between the total information content of the lognormal distribution and its correlation hierarchy (if the cutoff is at a large enough value, from (5.1) the two distributions have the same total amount of information). The hierarchy is thus still not well suited for the analysis of data in this regime.

For the same reason, even though any lognormal field is indeterminate, this effect plays no role for parameter inference in the linear regime, when the actual range of the variables is still small, and the tail at infinity is not yet felt. This is because in this regime on one hand the lognormal is still very close to a Gaussian over the range where it takes substantial values, and thus the lowest order correlations will still contain most of the

Fisher information, and on the other hand a few higher order terms are able to reproduce deviations of the functions $\partial_\alpha \ln p$ from the Gaussian very accurately over this small range. This is consistent with the findings in section 6.3.2 and 6.3.3 showing that the families presented there are indistinguishable from the lognormal for any practical purposes in the linear regime.

Let us comment in light of the criterion (6.27) on typical perturbative approaches in cosmology to parametrize (weakly) non Gaussian distributions. These involves moments, such as Gram-Charlier, Edgeworth expansions, or the relation between the moment generating function and the distribution (Bernardeau, 1994; Bernardeau and Kofman, 1995; Blinnikov and Moessner, 1998; Colombi, 1994; Fry, 1985; Juszkiewicz et al., 1995, e.g.), in one or several dimensions. It is therefore interesting to see to what extent they fit into this picture. Typically, when applied to the δ field, to first order these parametrize the non-Gaussianity through a polynomial with coefficients involving the cumulants, or equivalently the moments of the variable. Schematically,

$$p_\nu(\nu) \propto e^{-\nu^2/2} \left(1 + \alpha_3 H_3(\nu) + \alpha_4 H_4(\nu) + \cdots\right), \tag{6.28}$$

with $\nu = \delta/\sigma_\delta$. The coefficient α_i depends on the first i moments. The correction is given in terms of Hermite polynomials H_n, which are the orthogonal polynomials associated to the Gaussian distribution. Such expansions never produce a tailed distribution, in the sense that (6.27) is always met. The decay of the distribution namely still is Gaussian. Now, to first order and over the range of p, equation (6.28) is equivalent to

$$\ln p_\nu(\nu) \approx \text{cst} - \nu^2/2 + \alpha_3 H_3(\nu) + \alpha_4 H_4(\nu) + \cdots \tag{6.29}$$

Therefore, the functions $\partial_\alpha \ln p$ will have close to polynomial form. This is perfectly consistent with that decomposition of the Fisher information. Indeed, this expansion creates a probability density for which its Fisher information content is within the moments that were used to build it. This is another way to see that moment-indeterminate distributions cannot be produced by perturbative expansions

We now present both continuous as well as discrete families of probability density functions that have the same correlations as the lognormal at all orders, for any dimensionality d. In fact, it turns out that a stronger statement is true : for these families, all observables of the form

$$\langle \rho(x_1)^{n_1} \cdots \rho(x_d)^{n_d} \rangle, \quad n_i = \cdots -1, 0, 1 \cdots \tag{6.30}$$

are identical to those of the lognormal field, i.e. any power n_i can also be negative as well. Including the hierarchy of inverse powers and 'mixed' powers to the usual hierarchy thus still does not provide a complete description.

These families are generalizations to any number of dimension, means and two-point correlations of known one dimensional examples that can be found in the statistical literature (Heyde, 1963; Stoyanov, 1987).

Requirements such as homogeneity and isotropy are actually not needed in the following section. In particular, unless otherwise specified, $\bar{A}$ is a d-dimensional mean vector

$(\bar{A}(x_1), \cdots, \bar{A}(x_d))$, whose components can differ in principle. Nevertheless, the picture we have in mind is that of statistically homogeneous isotropic fields in a box of volume V, where some set of Fourier modes $k_{\min}$ to $k_{\max}$ can be probed. The corresponding Fourier representation of the two point correlations, in a continuous notation, is

$$[\xi_{A,\delta}]_{ij} = \int \frac{d^3k}{(2\pi)^3} P_{A,\delta}(k) e^{ik\cdot(x_i - x_j)} = \xi_{A,\delta}(x_i - x_j), \tag{6.31}$$

where the integral runs over these modes, and $\xi_{A,\delta}(r)$ is the ordinary two-point correlation function of A or δ. The matrix inverse is given by

$$\left[\xi^{-1}_{A,\delta}\right]_{ij} = \int \frac{d^3k}{(2\pi)^3} \frac{1}{P_{A,\delta}(k)} e^{ik\cdot(x_i - x_j)}. \tag{6.32}$$

This representation allow us to define a bit more rigorously what we mean by linear and non-linear lognormal field, or linear and non-linear regime, in the following discussion : if needed, it can be formally set as $P_A(k) \to 0$ or $P_A(k) \to \infty$ respectively, for all k.

6.3.2 Continuous family

Define the statistics of $\rho = (\rho(x_1), \cdots, \rho(x_d))$ through the following. Pick a real number ϵ with $|\epsilon| \leq 1$. Pick further a set of angular frequencies $\boldsymbol{\omega} = (\omega_1, \cdots, \omega_d)$. Each of these must be an integer. Fix $p^{LN}_{\rho}(\rho)$ the d-dimensional lognormal distribution with mean $\bar{A}$ and covariance matrix ξ_A defined earlier. Then set

$$p_\rho(\rho) := p^{LN}_{\rho}(\rho) \left[1 + \epsilon \sin\left(\pi \boldsymbol{\omega} \cdot \xi_A^{-1}\left(A - \bar{A}\right)\right)\right] \tag{6.33}$$

Since $|\epsilon| \leq 1$ this is positive and seen to be a well defined probability density function[1]. The claim that $p_\rho(\rho)$ defined in this way has the same moments $m_{\mathbf{n}}$ as the lognormal for any multiindex $\mathbf{n}$ is proved in the appendix. Note that in the above definition, $\bar{A}$ is the quantity that enters the definition of lognormal variables in equation (6.6). It is however not the mean of $A = \ln \rho$ anymore, when ρ is defined through (6.33).

The functional form of $p_\rho(\rho)$ consists of the lognormal envelope modulated by sinusoidal oscillations in A. The smaller the two-point function the higher frequency the oscillations. This may sound curious at first, since it seems to imply that the more linear the field, the more different the distributions within this family will thus appear. However, this is precisely when the oscillations are the strongest that this effect is less relevant. This can be seen as the following. Taking the average of any function f with respect to p_ρ leads trivially to

$$\langle f \rangle = \langle f \rangle_{LN} + \epsilon \left\langle f \sin\left(\pi \boldsymbol{\omega} \cdot \xi_A^{-1}\left(A - \bar{A}\right)\right)\right\rangle_{LN}, \tag{6.34}$$

where the subscript $_{LN}$ denotes the average with respect to the lognormal distribution. In the limit of the very linear regime, other terms fixed, the second term will average

[1] For $d = 1$, there are very slight differences with Heyde original family. Heyde unnecessarily writes 2π instead of π, and restricts ϵ and ω to be positive.

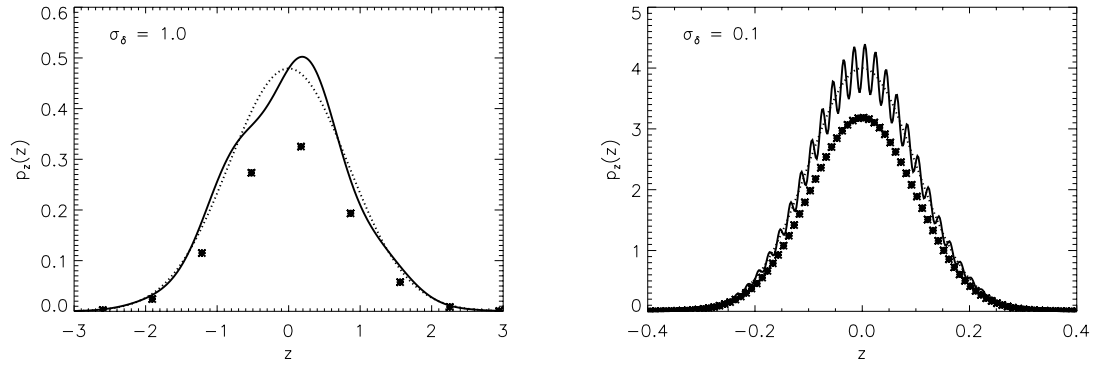

Fig. 6.2 — Left panel : Three different one dimensional distributions for $z := A - \bar{A}$, with identical moments $\langle \rho^n \rangle$, $\rho = e^A$, for all integer n, positive or negative. The dashed line is the zero mean Gaussian distribution, so that ρ is lognormal. The solid the member of the family in (6.33) with the lowest possible frequency, and amplitude $\epsilon = 0.1$. The discrete one is (6.39) with shift parameter $\alpha = 0.25$. They are shown at the scale of non linearity $\sigma_\delta = 1$, where this indeterminacy starts to become very relevant for inference. The families in any dimension are qualitatively identical to these. Right panel : Same as the left panel with $\sigma_\delta = 0.1$, when the indeterminacy if far less relevant for inference, for the reasons given in the text. The discrete distribution has been scaled by a constant factor for convenience.

out to zero for any reasonable f, since it is the integral of an highly oscillating function weighted by a smooth integrand. In the non-linear regime this in general ceases to be the case. This is illustrated as the solid lines in the left ($\sigma_\delta = 1$) and right ($\sigma_\delta = 0.1$) panels of figure 6.2, showing the member of that family in one dimension with minimal frequency $\omega = 1$, and $\epsilon = 0.1$. The dotted lines on these figures are the usual Gaussian for $A - \bar{A} = z$. The probability density function for $\ln \rho$ is not purely Gaussian anymore. It is therefore of interest to see how the correlations of A deviate from those of Gaussian variables. For instance the means $\left\langle A - \bar{A} \right\rangle$ do not vanish anymore as for the lognormal. A straightforward calculation leads to

$$\left\langle (A - \bar{A})(x_i) \right\rangle = -\epsilon \, \pi \omega_i \exp\left(-\frac{\pi^2}{2} \boldsymbol{\omega} \cdot \xi_A^{-1} \boldsymbol{\omega} \right). \tag{6.35}$$

Picking $\boldsymbol{\omega}$ as having a single non zero entry, ω, at x_i we get that they can be as large as

$$\begin{aligned} \left\langle (A - \bar{A})(x_i) \right\rangle &= -\epsilon \, \pi \omega \exp\left(-\frac{\pi^2}{2} w^2 \left[\xi_A^{-1} \right]_{ii} \right) \\ &:= -\epsilon \, \pi \omega \exp\left(-\frac{\pi^2}{2} \frac{\omega^2}{\sigma^2_{A,\mathrm{eff}}(x_i)} \right) \end{aligned} \tag{6.36}$$

Observables as simple as the means of A are therefore not constrained by the knowledge of the entire correlation hierarchy of the lognormal field. While the effect is irrelevant in the linear regime (for say $\sigma_{A,\mathrm{eff}} = 0.1$, the maximal value of the mean in equation

(6.36) is only $\approx 10^{-215}$), deep in the non-linear regime this is not the case anymore. It is easy to show from the above expression that the range available to $\left\langle \left(A - \bar{A}\right)(x_i) \right\rangle$, choosing ω appropriately, scales to infinity with $\propto \sigma_{A,\text{eff}}$. The means are thus left totally unconstrained in that regime. This and the very sharp behavior is of course a generic effect, not limited to that particular observable. It is obvious that the relevance of this effect for parameter inference is very sensitive to the degree of linearity of the field, and that large amounts of information are lost to the hierarchy in the high variance regime[2].

6.3.3 Discrete family

Fix again the dimensionality d, the vector $\bar{A}$ and the matrix ξ_A. For all integer valued d-dimensional multiindex $\mathbf{n}$ define a realization $A_\mathbf{n}$ of A as the following. Pick $\alpha = (\alpha_1, \cdots, \alpha_d)$ any point, and set

$$A_\mathbf{n} := \bar{A} + \xi_A \cdot (\mathbf{n} - \alpha)\,. \tag{6.37}$$

While α can in principle be anything, only components $\alpha_i \in [0, 1)$ will actually define different grids. As usual, ρ is given by exponentiation,

$$\rho_\mathbf{n} := \exp\left(A_\mathbf{n}\right) \tag{6.38}$$

Assign then to these realizations parametrized by $\mathbf{n}$ a probability

$$P_\mathbf{n} = \frac{1}{Z} \exp\left(-\frac{1}{2}\left(A_\mathbf{n} - \bar{A}\right) \cdot \xi_A^{-1} \left(A_\mathbf{n} - \bar{A}\right)\right). \tag{6.39}$$

These are usual Gaussian probabilities for $A_\mathbf{n}$, except that we have only a discrete set of field realizations. Note that it can be written, maybe more conveniently, as

$$P_\mathbf{n} = \frac{1}{Z} \exp\left(-\frac{1}{2}\left(\mathbf{n} - \alpha\right) \cdot \xi_A \left(\mathbf{n} - \alpha\right)\right). \tag{6.40}$$

Since ξ_A is positive definite, the normalization factor Z is seen to be well defined, as for more usual Gaussian integrals, and so are the probabilities. This discrete probability distribution has the same moments of $\rho_\mathbf{n}$ than the d-dimensional lognormal distribution with associated $\bar{A}$ and ξ_A, as proven in the appendix 6.3.4. Again, negative entries in $\mathbf{n}$ are allowed.

This family is clearly different from the previous, continuous one. Rather than modulating the lognormal distribution with an oscillating factor, it is a series of Dirac delta functions sampling the lognormal on the grid given by (6.37). The role of α is to shift the sample by a small amount. If α is set to zero, then $A = \bar{A}$ is part of the sample, while it is not if not. The fact that this indeterminacy is irrelevant in the linear regime comes this

[2] Among this family, it turns out that some observables such as the variances $\left\langle (A - \bar{A})^2(x_i) \right\rangle$ are always identical to $\sigma_A^2(x_i)$ for any choice of ϵ and $\boldsymbol{\omega}$. We do not attach any significance to this, since this is not the case for the discrete family, though closed analytical expressions cannot be obtained in this case.

time from realizing that for any nice enough function f, the average of f will converge to $\langle f \rangle_{LN}$ due to the trapezoidal rule of quadrature. The grid spacing at which A is sampled in this way in (6.38) becomes namely thinner and thinner. In the non-linear regime, the spacing is however very large, leading again to large deviations. This is also illustrated in figure 6.2 for the one dimensional version of it, with shift parameter $\alpha = 0.25$.

6.3.4 Appendix

We prove the claim that the distributions we defined have the same correlations than the lognormal at all orders. As we will see this is also true including 'negative orders' and 'mixed orders', i.e. when negative powers of the variables are allowed in the correlations. Recall that for lognormal variables $\rho = (\rho_1, \cdots, \rho_d)$ with means and covariance matrix of their logarithms $\bar{A} = (\bar{A}_1, \cdots \bar{A}_d)$ and ξ_A we have

$$m_{\mathbf{n}} := \langle \rho^{\mathbf{m}} \rangle = \left\langle \rho_1^{n_1} \cdots \rho_d^{n_d} \right\rangle = \exp\left(\mathbf{n} \cdot \bar{A} + \frac{1}{2} \mathbf{n} \cdot \xi_A \mathbf{n} \right), \quad \mathbf{n} = (n_1, \cdots, n_d). \tag{6.41}$$

A simple proof of this fact is to make use of the standard formulae for Gaussian integrals, valid for any positive matrix ξ_A, mean vector A and vector z, that can be complex valued.

$$\frac{1}{(2\pi)^{d/2}} \frac{1}{\sqrt{\det \xi_A}} \int d^d A \exp\left(-\frac{1}{2} \left(A - \bar{A}\right) \cdot \xi_A^{-1} \left(A - \bar{A}\right) + \left(A - \bar{A}\right) \cdot z \right) = \exp\left(\frac{1}{2} z \cdot \xi_A z \right). \tag{6.42}$$

Essentially all calculations in this work follow from this formula. Even the proof for the discrete family can be considered a discrete version of that relation.

Continuous family To prove our claim it is enough to show that

$$\left\langle \rho^{\mathbf{n}} \sin\left(\pi\, \boldsymbol{\omega} \cdot \xi_A^{-1} \left(A - \bar{A}\right) \right) \right\rangle_{LN} = 0. \tag{6.43}$$

This must hold for any d-dimensional multiindices $\boldsymbol{\omega}$ and $\mathbf{n}$ (we allow entries to be negative), where the average is taken with respect to the lognormal density function, equation (6.8). We proceed as the following : we evaluate the following integral

$$I(\mathbf{n}, \boldsymbol{\omega}) := \left\langle \rho^{\mathbf{n}} \exp\left(i\pi\, \boldsymbol{\omega} \cdot \xi_A^{-1} \left(A - \bar{A}\right) \right) \right\rangle_{LN}, \tag{6.44}$$

and show that its imaginary part vanishes for $\boldsymbol{\omega}$ and $\mathbf{n}$ as specified.
Writing equation (6.44) using

$$\rho^{\mathbf{n}} = \exp(\mathbf{n} \cdot A) = \exp\left[\mathbf{n} \cdot \left(A - \bar{A}\right) + \mathbf{n} \cdot \bar{A} \right] \tag{6.45}$$

leads immediately to the Gaussian integral given in (6.42), with $z = \mathbf{n} + i\pi \xi_A^{-1} \boldsymbol{\omega}$. It follows from that equation

$$I(\mathbf{n}, \boldsymbol{\omega}) = \exp\left[\mathbf{n} \cdot \bar{A} + \frac{1}{2} \left(\mathbf{n} + i\pi \xi_A^{-1} \boldsymbol{\omega}\right) \cdot \xi_A \left(\mathbf{n} + i\pi \xi_A^{-1} \boldsymbol{\omega}\right) \right]. \tag{6.46}$$

Separating real from imaginary argument, this expression reduces to

$$I(\mathbf{n},\boldsymbol{\omega}) = \exp\left(\mathbf{n}\cdot\bar{A} + \frac{1}{2}\mathbf{n}\cdot\xi_A\mathbf{n} - \frac{\pi^2}{2}\boldsymbol{\omega}\cdot\xi_A^{-1}\boldsymbol{\omega}\right)\cdot\exp\left(i\pi\,\boldsymbol{\omega}\cdot\mathbf{n}\right). \tag{6.47}$$

The imaginary part of that expression is thus proportional to $\sin\pi\,\boldsymbol{\omega}\cdot\mathbf{n}$. Whenever $\boldsymbol{\omega}$ and $\mathbf{n}$ are integer valued, so is their scalar product $\boldsymbol{\omega}\cdot\mathbf{n} = \sum_i \omega_i n_i$. Therefore, the sine vanishes and (6.43) is proved.

Discrete family From equation (6.37) and (6.38), we have

$$\rho_{\mathbf{n}}^{\mathbf{m}} = \exp\left(\mathbf{m}\cdot\bar{A} + \mathbf{m}\cdot\xi_A\left(\mathbf{n}-\alpha\right)\right). \tag{6.48}$$

It follows that the moments of ρ are given by

$$\langle\rho^{\mathbf{m}}\rangle = \frac{e^{\mathbf{m}\cdot\bar{A}}}{Z}\sum_{\mathbf{n}\in\mathbb{Z}^d}\exp\left[-\frac{1}{2}\left(\mathbf{n}-\alpha\right)\xi_A\left(\mathbf{n}-\alpha\right) + \mathbf{m}\cdot\xi_A\left(\mathbf{n}-\alpha\right)\right]. \tag{6.49}$$

The proof is based on completing the square in the exponent, in perfect analogy of standard proofs of the Gaussian integral in (6.42). Write

$$-\frac{1}{2}\left(\mathbf{n}-\alpha\right)\cdot\xi_A\left(\mathbf{n}-\alpha\right) + \mathbf{m}\cdot\xi_A\left(\mathbf{n}-\alpha\right) = -\frac{1}{2}\left(\mathbf{n}-\mathbf{m}-\alpha\right)\xi_A\left(\mathbf{n}-\mathbf{m}-\alpha\right) + \frac{1}{2}\mathbf{m}\cdot\xi_A\mathbf{m}, \tag{6.50}$$

and then perform the shift of summing index $\mathbf{n}\to\mathbf{n}+\mathbf{m}$, obtaining

$$\langle\rho^{\mathbf{m}}\rangle = \exp\left(\mathbf{m}\cdot\bar{A} + \frac{1}{2}\mathbf{m}\cdot\xi_A\mathbf{m}\right)\frac{1}{Z}\sum_{\mathbf{n}\in\mathbb{Z}^d}\exp\left(-\frac{1}{2}\left(\mathbf{n}-\alpha\right)\cdot\xi_A\left(\mathbf{n}-\alpha\right)\right). \tag{6.51}$$

Since the sum ranges over all the multiindices, the shift does not create boundary terms. This last sum is nothing else than Z, so that we recover

$$\langle\rho^{\mathbf{m}}\rangle = \exp\left(\mathbf{m}\cdot\bar{A} + \frac{1}{2}\mathbf{m}\cdot\xi_A\mathbf{m}\right), \tag{6.52}$$

which are indeed the same as the lognormal in (6.41). Again, this is also true if negative entries in $\mathbf{m}$ are permitted.

6.4 Information at all orders for the one dimensional distribution

We now obtain the information coefficients s_n, and thus the matrices $F_{\leq N}$, for the lognormal distribution in one dimension $d = 1$ at all orders. By additivity of the information order by order this is equivalent to obtain that of the uncorrelated lognormal field in any number of dimensions.

There are two free parameters, the mean $\bar{A}$ and the variance σ_A^2. The dependency on cosmological parameters $\boldsymbol{\theta} = (\alpha, \beta, \cdots)$ can enter one or both of these parameters. We have from section 6.1.2

$$\sigma_A^2 = \ln\left(1+\sigma_\delta^2\right) \quad \bar{A} = \ln\bar{\rho} - \frac{1}{2}\ln\left(1+\sigma_\delta^2\right), \tag{6.53}$$

where σ_δ^2 is the variance of the fluctuations.

Note that the total Fisher information content of ρ becomes

$$F_{\alpha\beta} = \frac{1}{\sigma_A^2}\frac{\partial\bar{A}}{\partial\alpha}\frac{\partial\bar{A}}{\partial\beta} + \frac{1}{2\sigma_A^4}\frac{\partial\sigma_A^2}{\partial\alpha}\frac{\partial\sigma_A^2}{\partial\beta}. \tag{6.54}$$

The key parameter throughout this section will be the quantity q, defined as

$$q := e^{-\sigma_A^2} = \frac{1}{1+\sigma_\delta^2}. \tag{6.55}$$

Note that q is strictly positive and smaller than unity. The regime of small fluctuations, where the lognormal distribution is very close to the Gaussian distribution is described by values of q close to unity. Deep in the non linear regime, it tends to zero. These two regimes are conveniently separated at $q = 1/2$, corresponding to fluctuations of unit variance.

The moments of the distribution are given by

$$m_n = \bar{\rho}^n q^{-\frac{1}{2}n(n-1)}. \tag{6.56}$$

We note the following extremely convenient property of these moments,

$$m_{i+j} = m_i m_j \, q^{-ij}. \tag{6.57}$$

From chapter 3, equation (3.8) we see that the nth information coefficient s_n is given by

$$s_n(\alpha) = \frac{\partial\ln\bar{\rho}}{\partial\alpha}\sum_{k=0}^{n} C_{nk}\, m_k\, k + \frac{\partial\sigma_\delta^2}{\partial\alpha}\frac{q}{2}\sum_{k=0}^{n} C_{nk}\, m_k\, k(k-1), \tag{6.58}$$

where the nth orthornormal polynomial P_n is given by $P_n(\rho) = \sum_{k\leq n} C_{nk}\rho^k$, with yet unknown matrix elements C_{nk} to be found.

Evaluation of the above sums can proceed in different ways. Notably, it is possible to get an explicit formula for the orthonormal polynomials, and therefore of the matrix C, for the lognormal distribution. These are essentially the Stieltjes-Wigert polynomials (Szegö, 2003; Wigert, 1923). We will namely use their specific form later in this section, though they are not needed for the purpose of evaluating (6.58).

We proceed with the following trick : we introduce the q-shifted factorial , also called q-Pochammer symbol (Andrews et al., 1999; Kac and Cheung, 2001, section 10), as

$$(t:q)_n := \prod_{k=0}^{n-1} \left(1 - tq^k\right), \quad (t:q)_0 := 1 \tag{6.59}$$

t a real number, and prove (see appendix 6.4.4 for the proof) that the following curious identity holds,

$$\langle P_n(t\rho)\rangle = (-1)^n \frac{q^{n/2}}{\sqrt{(q:q)_n}} (t:q)_n \,. \tag{6.60}$$

By virtue of

$$\langle P_n(t\rho)\rangle = \sum_{k=0}^{n} C_{nk}\, m_k\, t^k, \tag{6.61}$$

it follows from our identity (6.60) that the sums given in the right hand side of equation (6.58) are proportional to the first, respectively the second derivative of the q-Pochammer symbol evaluated at $t = 1$. Besides, matching the powers of t on both sides of equation (6.60) will provide us immediately the explicit expression for the matrix elements C_{nk}.

We consider now the two cases of the parameter of interest α being $\alpha = \ln\bar{\rho}$ and $\alpha = \sigma_\delta^2$. The general case of a generic parameter is reconstructed trivially from these two from the chain rule of derivation. The total amount of information in the distribution becomes in the first case from (6.54)

$$I(\ln\bar{\rho}) = -\frac{1}{\ln q}. \tag{6.62}$$

The second case is the most common in cosmology, for instance for any model parameter entering the matter power spectrum. The exact amount of information on this parameter is this time

$$I(\sigma_\delta^2) = -\frac{q^2}{\ln q}\left(\frac{1}{4} - \frac{1}{2\ln q}\right). \tag{6.63}$$

In both of these situations, we obtain the information coefficients (6.58) by differentiating once, respectively twice, our relation (6.60) with respect to the parameter t, and evaluating these derivatives at $t = 1$. The result is

$$s_n(\ln\bar{\rho}) = (-1)^{n-1}\sqrt{\frac{q^n}{1-q^n}\,(q:q)_{n-1}} \tag{6.64}$$

and

$$s_n(\sigma_\delta^2) = -s_n(\ln\bar{\rho})q\left[\sum_{k=1}^{n-1}\frac{q^k}{1-q^k}\right], \quad n > 1 \tag{6.65}$$

whereas $s_{n=1}(\sigma_\delta^2)$ is easily seen to vanish, since the mean is independent of σ_δ^2.

6.4.1 Lack of information in the moments

The series

$$i(\alpha) := \sum_{n=1}^{\infty} s_n^2(\alpha) \tag{6.66}$$

are the total amount of information contained in the full series of moments. The efficiencies $\epsilon(\alpha)$ defined as

$$\epsilon(\alpha) := \frac{i(\alpha)}{I(\alpha)}. \tag{6.67}$$

are the fraction of the information that can be accessed by extraction of the full set of moments of our lognormal variable. Similarly, we define

$$\epsilon_N(\alpha) := \frac{1}{I(\alpha)} \sum_{n=1}^{N} s_n^2(\alpha) \tag{6.68}$$

for the first N moments. Note that from the chain rule of derivation, both $\epsilon(\alpha)$ and $\epsilon_N(\alpha)$ are invariant under a reparametrisation of α, i.e.

$$\epsilon(\beta) = \epsilon(\alpha) \text{ whenever } \alpha = f(\beta) \tag{6.69}$$

for some function f, all other parameters fixed. For this reason, these efficiencies refer rather to a parameter type (in our case mean alike or variance alike) rather than some more specific instance.

The two asymptotic regimes of very small and very large fluctuation variance σ_δ can be seen without difficulty. Both for $\alpha = \ln \bar{\rho}$ and $\alpha = \sigma_\delta^2$, it is seen that in these asymptotic regimes the first non vanishing term of the corresponding series (6.64) and (6.65) dominates completely its value. For very small variance, or equivalently q very close to unity, both efficiencies ϵ tend to unity, illustrating the fact the distribution becomes arbitrary close to Gaussian : all the information is contained in the first two moments. The large variance regime is more interesting, and, even tough the information coefficients decays very sharply as well, the series (6.66) are far from converging to the corresponding expressions (6.62) and (6.63) showing the total amount of information. Considering only the dominant first term in the relevant series and setting $q \to 0$, one obtains

$$\epsilon(\ln \bar{\rho}) \to \frac{1}{\sigma_\delta^2} \ln\left(1 + \sigma_\delta^2\right). \tag{6.70}$$

and a much more dramatic decay of $\epsilon(\sigma_\delta^2)$:

$$\epsilon(\sigma_\delta^2) \to \frac{4}{\sigma_\delta^8} \ln\left(1 + \sigma_\delta^2\right). \tag{6.71}$$

Both series given in (6.66) are quickly (exponentially) convergent and well suited for numerical evaluation. Figure 6.3 shows the accessible fractions ϵ of information through extraction of the full series moments. Figure 6.4 shows the repartition of this accessible

fraction among the first 10 moments. Most relevant from a cosmological point of view in figure 6.3 is the solid line, dealing with the case of the parameters of interest entering the variance only. These figures shows clearly that the moments, as probes of the lognormal matter field, are penalized by two different processes. First, as soon as the field shows non-linear features, following equations (6.70) and (6.71), almost the entirety of the information content cannot be accessed anymore by extracting its successive moments. Within a range of one magnitude in the variance, the moments goes from very efficient probes to highly inefficient. Second, as shown in figure 6.4, in the still close to linear regime, as the variance of the field approaches unity, this accessible fraction gets quickly transferred from the variance alone to higher order moments.

This repartition of the information within the moments is built out of two different regimes. First, for large variance, or large n, we see easily from the above expressions (6.64) and (6.65) that in both cases the information coefficients decays exponentially,

$$s_n^2 \propto q^n, \quad -n \ln q \gg 1. \tag{6.72}$$

On the other hand, if the variance or n is small enough, we can set $1 - q^n \approx -n \ln q$, and we obtain, very roughly,

$$s_n^2 \propto [-n \ln q]^n, \quad -n \ln q \ll 1, \tag{6.73}$$

explaining the trend with variance seen in figure 6.4, that puts more importance to higher order moments as the variance grows. Note that the latter regime can occur only for small enough values of the variance. Deeper in the non linear regime, the trend is therefore reversed, obeying (6.72) for all values of n, with a steeper decay for higher variance. This is clearly shown in the right panel of figure 6.4. In that regime, higher order moments do not carry additional information.

6.4.2 A q-analog of the logarithm

These results show clearly that large parts of the information become invisible to the moments. However, it does not tell us what is responsible for this phenomenon. It is therefore of interest to look into more details of these missing pieces of information. As we have seen, these are due to the inability of the polynomials to reconstruct precisely the score function. In the case $\alpha = \ln \bar{\rho}$, the score function $\partial_\alpha \ln p$ of our lognormal distribution is easily shown to take the form of a logarithm in base q,

$$s(\rho, \ln \bar{\rho}) = -\frac{1}{2} - \ln_q \left(\frac{\rho}{\bar{\rho}} \right). \tag{6.74}$$

Therefore the series

$$s_\infty(\rho, \ln \bar{\rho}) := \sum_{n=0}^{\infty} s_n(\ln \bar{\rho}) P_n(\rho) \tag{6.75}$$

will represent some function, very close to the logarithm (6.74) for $q \to 1$ over the range of $p(\rho, \boldsymbol{\theta})$. It will however fail to reproduce some of its features at lower q-values. This

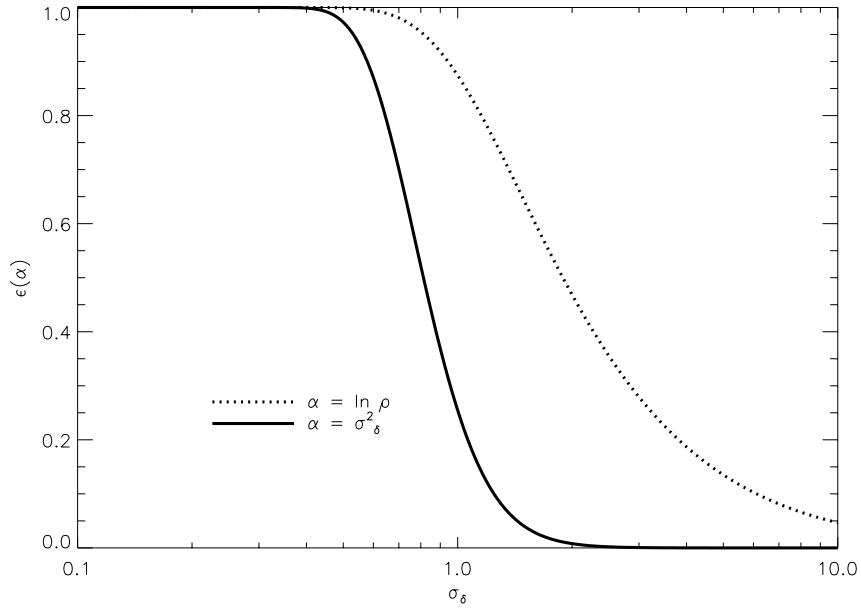

Fig. 6.3 — The fraction of the total information content that is accessible through extraction of the full series of moments of the lognormal field, as function of the square root of the variance of the fluctuations. The solid line for a parameter entering the variance of the field and not its mean, and dotted conversely.

is hardly surprising, since it is well known that the logarithm function does not have a

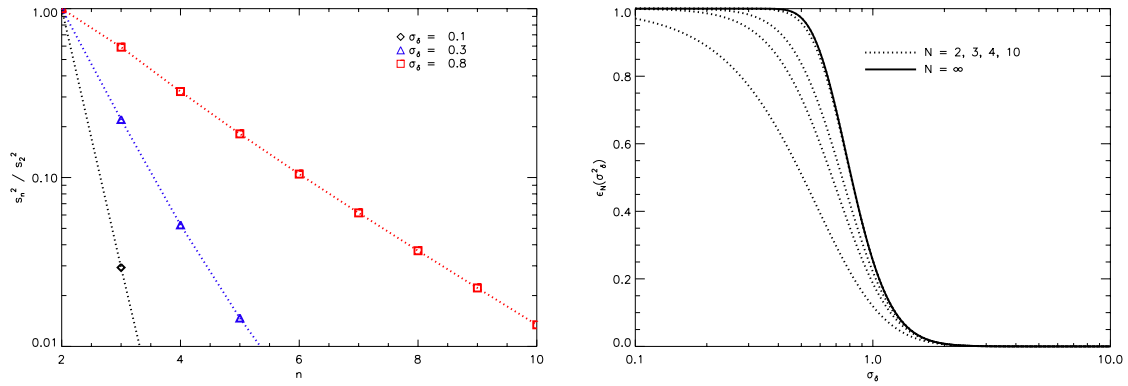

Fig. 6.4 — Left panel : The distribution of the information within the first 10 moments of the lognormal field, given by the coefficients $s_n^2(\sigma_\delta^2)$, equation (6.65), normalized to the information content of the second moment, for three different values of σ_δ. Note that deeper in the non linear regime, the trend is reversed. Right panel : the cumulative efficiencies of the moments to capture the information on σ_δ^2, as function of the root of the fluctuations.

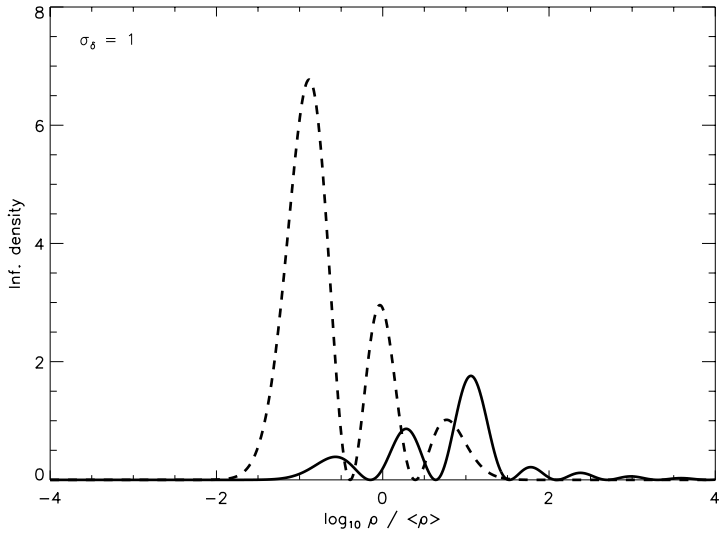

Fig. 6.5 — The information density of the lognormal distribution, dashed, and, solid, its approximation through the associated orthonormal polynomials, for a variance alike parameter, for fluctuations of unit variance. While most of the information of the lognormal field in this regime is actually contained in the underdense regions, the moments are essentially unable to catch it.

Taylor expansion over the full positive axis. For this reason, the approximation

$$s_\infty(\rho, \ln\bar{\rho}) := \sum_{n=0}^{\infty} s_n(\ln\bar{\rho}) P_n(\rho) \tag{6.76}$$

of the score function $s(\rho, \ln\bar{\rho})$ through polynomials can indeed only fail when the fluctuation variance becomes large enough. In the appendix, we show that $s_\infty(\rho, \ln\bar{\rho})$ takes the form

$$s_\infty(\rho, \ln\bar{\rho}) = -\sum_{k=1}^{\infty} \frac{q^k}{1-q^k} \left[1 + (-1)^k \frac{q^{k(k-1)}}{(q:q)_k} \left(\frac{\rho}{\bar{\rho}}\right)^k\right]. \tag{6.77}$$

It is interesting to note that this series expansion is almost identical to the one of the q-analog of the logarithm S_q defined by E. Koelink and W. Van Assche, with the only difference being the replacement of $q^{k(k-1)/2}$ by $q^{k(k-1)}$ (See Koelink and Van Assche (2009), and also Gautschi (2008)). Due to this replacement, s_∞ does not possess several properties S_q has and makes it a real q-analog of the logarithm, such as $S_q(q^{-n}) = n$, for positive integers. The qualitative behavior of s_∞ stays however close to S_q. Notably, its behavior in underdense regions, $\rho/\bar{\rho} \ll 1$, where as seen from (6.77) s_∞ tends to a finite value, is very different from a logarithm.

This calculation can be performed as well in the case $\alpha = \sigma_\delta^2$, with similar conclusions. Since it is rather tedious and not very enlightening, we do not reproduce it in these pages. We show in figure 6.5 the information density of the lognormal distribution (dashed line),

and its approximation by the orthogonal polynomials (solid line),

$$p(\rho, \boldsymbol{\theta}) \left(\sum_{n=0}^{\infty} s_n(\sigma_\delta^2) P_n(\rho) \right)^2 , \tag{6.78}$$

when the fluctuation variance σ_δ^2 is equal to unity. It is clear from this figure that in this regime, while most of information is located within the underdense regions of the lognormal field, the moments are however unable to catch it.

As a non trivial check of the correctness of our numerical and analytical calculations, we compared the total information content as evaluated from integrating the information densities on figure 6.5 to the one given by the equation (6.63), respectively (6.66), with essentially perfect agreement.

6.4.3 Comparison to standard perturbation theory

As we have seen in chapter 3, the knowledge of its first $2n$ moments allows for any distribution the direct evaluation of the independent information content of the first n moments, for instance from equation (3.14). This even if the exact shape of the distribution is not known, or too complicated. In particular, we can use the explicit expressions for the first six moments of the density fluctuation field within the framework of standard perturbation theory (SPT) provided by F. Bernardeau in (Bernardeau, 1994), in order to compare $s_2(\alpha)$ and $s_3(\alpha)$ as given from SPT to their lognormal analogs.

We note that a comparison to (Bernardeau, 1994) can only be very incomplete and, to some extent, it can only fail. It is indeed part of the approach in (Bernardeau, 1994), when producing functional forms for the distribution of the fluctuation field, to invert the relation between a moment generating function and its probability density function. For such an inversion to be possible it is of course necessary that the probability density is uniquely determined by its moments. As said, this is not the case for the lognormal distribution. Therefore, that approach can never lead to an exact lognormal distribution, or to any distribution for which the moment hierarchy forms an incomplete set of probes. However, such a comparison can still lead to conclusions relevant for many practical purposes, such as those dealing with the first few moments.

The variance of the field is explicitly given as an integral over the matter power spectrum,

$$\sigma_\delta^2 = \frac{1}{2\pi^2} \int_0^\infty dk\, k^2 P(k, \boldsymbol{\theta}) \left| W(kR) \right|^2 , \tag{6.79}$$

where $W(kR)$ is the Fourier transform of the real space top hat filter of size R, and any cosmological parameter α entering the power spectrum $P(k)$. In the notation of (Bernardeau, 1994), the moments of the fluctuation field $m_n = \langle \delta^n \rangle$ are given by the deconnected, or Gaussian, components, while the connected components $\langle \delta^n \rangle_c$, $n \geq 3$

are given in terms of parameters S_n,

$$\langle \delta^n \rangle_c = \sigma^{2(n-1)} S_n. \tag{6.80}$$

The parameters S_n contain a leading, scale independent coefficient, and deviation from this scale independence are given in terms of the logarithmic derivative of the variance,

$$\gamma_i = \frac{d^i \ln \sigma_\delta^2}{d \ln R^i}, \quad i = 1, \cdots \tag{6.81}$$

Neglecting the very weak dependence of S_n on cosmology, from (6.80) we can write

$$\frac{\partial m_n}{\partial \alpha} = \frac{\partial \sigma_\delta^2}{\partial \alpha} \cdot \begin{cases} 0, & n = 1 \\ 1, & n = 2 \\ 2m_3 \,/\, \sigma_\delta^2, & n = 3 \end{cases} \tag{6.82}$$

With the coefficients S_n up to $n = 6$ given in (Bernardeau, 1994, page 703), and the above relations, we performed a straightforward evaluation of the information coefficients $s_2^2(\alpha)$ and $s_3^2(\alpha)$, from equation (3.14). The variance was obtained from (6.79) within a flat ΛCDM universe ($\Omega_\Lambda = 0.7, \Omega_m = 0.3, \Omega_b = 0.045, h = 0.7$), with power spectrum parameters ($\sigma_8 = 0.8, n = 1$) and we used the transfer function from Eisenstein and Hu (Eisenstein and Hu, 1998). The needed derivatives γ_i, $i = 1, \cdots, 4$ were obtained numerically through finite differences.

In figure 6.6, we show the ratio

$$\left(\frac{s_3(\alpha)}{s_2(\alpha)} \right)^2, \tag{6.83}$$

both for the lognormal distribution and the SPT predictions, for $\alpha = \sigma_\delta^2$. Note that this ratio is actually independent of the parameter α as long as it does not enter the mean of the field but only the power spectrum. It is the relative importance of the third moment with respect to the second, as function of the variance, This ratio is identically zero for a Gaussian distribution. The models stands in good agreement over many orders of magnitude. It is striking that both models consistently predict that a the entrance of the non-linear regime, this ratio takes a maximal value close to unity. Surely, the SPT curve for larger values of the variance is hard to interpret, since out of its domain of validity.

6.4.4 Appendix

Derivation of 6.60 To prove (6.60), we note that both sides of the equation are polynomials of degree n in t, and that the zeroes of the right hand side are given by

$$t = q^{-i}, \quad i = 0, \cdots, n-1. \tag{6.84}$$

We first show that the left hand side evaluated at these points does vanish as well, so that the two polynomials must be proportional. We then find the constant of proportionality

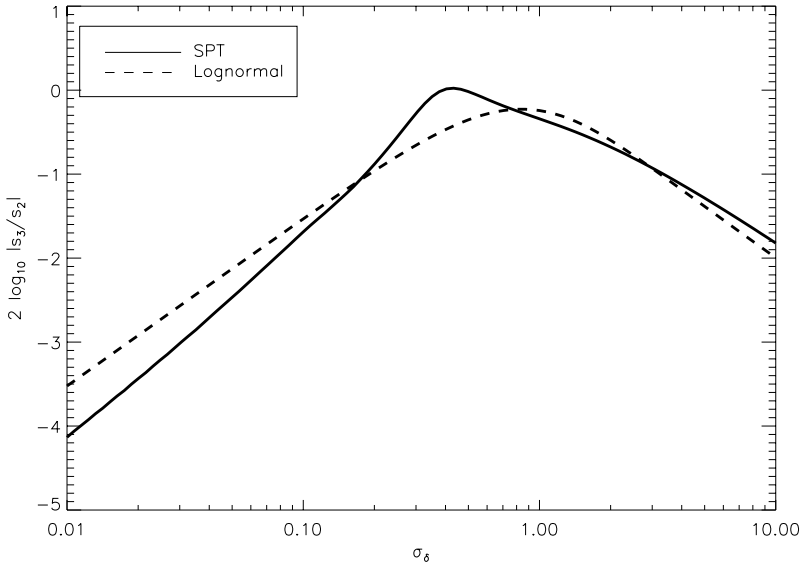

Fig. 6.6 — The ratio of the independent information content of the third moment to that of the second moment, for the lognormal field (dashed) and standard perturbation theory (solid), as function of the square root of the variance of the fluctuations.

by requiring P_n to have the correct normalization.

The first step is performed by noting that

$$\left\langle P_n(q^{-i}\rho)\right\rangle = \frac{1}{m_i}\left\langle P_n(\rho)\rho^i\right\rangle, \quad i = 0, 1, \cdots \tag{6.85}$$

an identity which is proven by expanding P_n in both sides of the equation in terms of the transition matrix C, and using the relation (6.57) between the moments. Since P_n is by construction orthogonal to any polynomial of strictly lower degree, we have indeed

$$\left\langle P_n(q^{-i}\rho)\right\rangle = 0, \quad i = 0, \cdots, n-1. \tag{6.86}$$

This implies

$$\sum_{k=0}^{n} C_{nk} m_k \, t^k = \alpha_n \left(t : q\right)_n \tag{6.87}$$

for some constant of proportionality α_n. To find it, we note that by expanding the normalization condition of P_n,

$$1 = \left\langle P_n^2(\rho)\right\rangle, \tag{6.88}$$

using again property (6.57), it must hold that

$$1 = \sum_{i,j=0}^{n} C_{ni} m_j \, C_{nj} m_j \, q^{-ij}. \tag{6.89}$$

The sums can be performed using equation (6.87), leading to the following equation for α_n,

$$1 = (-1)^n \, \alpha_n^2 \, q^{n(n-1)/2} \left(q^{-n} : q\right)_n . \tag{6.90}$$

This expression simplifies to

$$\alpha_n^2 = \frac{q^n}{(q:q)_n} \tag{6.91}$$

and the sign of α_n must be -1^n in order to have a positive matrix element C_{nn}. This concludes the proof of (6.60).

Derivation of the representation (6.77) In order to get the explicit series representation of (6.77), we first obtain from relation (6.60) the exact expression of the transition matrix C. The expansion of the q-Pochammer symbol on the right hand side of (6.60) in powers of t is the Cauchy binomial theorem,

$$(t:q)_n = \sum_{k=0}^{n} \binom{n}{k}_q q^{k(k-1)/2} (-t)^k, \tag{6.92}$$

where

$$\binom{n}{k}_q = \frac{(q:q)_n}{(q:q)_k \, (q:q)_{n-k}} \tag{6.93}$$

is the Gaussian binomial coefficient. Matching powers of t in (6.60) we obtain the explicit form

$$C_{nk} = (-1)^{n-k} \frac{q^{n/2}}{\sqrt{(q:q)_n}} \binom{n}{k}_q q^{k^2} \bar{\rho}^{-k}. \tag{6.94}$$

Therefore, interchanging the n and k sums in (6.77) , it holds

$$s_\infty(\rho, \ln \bar{\rho}) = -\sum_{n=1}^{\infty} \frac{q^n}{1-q^n} + \sum_{k=1}^{\infty} q^{k^2} \left(-\frac{\rho}{\bar{\rho}}\right)^k \sum_{n=k}^{\infty} \frac{q^n}{1-q^n} \binom{n}{k}_q . \tag{6.95}$$

With the help of some algebra the following identity is not difficult to show

$$\sum_{n=k}^{\infty} \frac{q^n}{1-q^n} \binom{n}{k}_q = \frac{1}{(q:q)_k} \frac{q^k}{1-q^k}, \quad k \geq 1. \tag{6.96}$$

Consequently, the series expansion of $s_\infty(\rho, \bar{\rho})$ is given by

$$s_\infty(\rho, \ln \bar{\rho}) = -\sum_{k=1}^{\infty} \frac{q^k}{1-q^k} \left[1 + (-1)^k \frac{q^{k(k-1)}}{(q:q)_k} \left(\frac{\rho}{\bar{\rho}}\right)^k\right] \tag{6.97}$$

6.5 Connection to N-body simulations

Mark C. Neyrinck analysed in (Neyrinck, 2011) the Coyote Universe N-body simulations suite (Heitmann et al., 2010; Lawrence et al., 2010) in a box of volume $V = 2.2\text{Gpc}^3$, with 256^3 cells, extracting the spectrum $P(k)$ of A and δ over the range $0.02/\text{Mpc} \lesssim k \lesssim 0.6/\text{Mpc}$, comparing their statistical power as function of the smallest scale $k_{\max}$ included in the analysis for several cosmological parameters. It was found that the spectrum of A has more constraining power on cosmological parameters than that of δ, when the non linear scales are included in the analysis. We refer to that paper for more details on the procedures and results. In this framework, ρ is $1+\delta$, and thus $A = \ln(1+\delta)$. The fields are statistically homogeneous and isotropic.

Given the considerations of the previous sections, and the fact that the density field is known to be somewhat close to lognormal, these results can hardly be considered surprising. The field A must be indeed closer to a Gaussian field for all values of the cosmological parameters, so that low order N point functions of A must contain a larger fraction of the information than those of δ (it is useful to remember that the full fields A and δ carry in all cases the very same total amount of information, since the mapping between them is parameter independent and invertible). In this section we want to go a step further from these qualitative considerations and make a quantitative comparison of these results to simple analytical methods using the results of the previous sections.

First, we need to make sure that a Gaussian description of the field A seen in the simulations is reasonable, at least for what concerns the information content. In particular, this is not the case for the smallest scales of A, since the covariance matrix of P_A in the 256^3 box clearly shows substantial off diagonal elements starting from $k \simeq 0.3/\text{Mpc}$. The same analysis was therefore repeated, performing the logarithmic transform on the δ field only after smoothing δ on twice the original length scale, by merging the 256^3 into 128^3 cells. This allowed us to extract the spectra of A and δ over the range $0.02/\text{Mpc} \lesssim k \lesssim 0.3/\text{Mpc}$, with a diagonal covariance matrix over the full range to a very good approximation. It is important to realize that sadly it is not identical to the much simpler approach of considering the original A field only up to the new $k_{\max}$: since all the scales of δ have an impact on the large scales of A, the operations of smoothing δ and then log transforming δ are not identical to log transforming δ and then smoothing A.

As we discussed already on several occasions, for a purely Gaussian field with spectrum P, the information content on α in the spectrum is given by

$$I(\alpha) = \frac{V}{2}\int \frac{d^3k}{(2\pi)^3}\left(\frac{\partial \ln P(k)}{\partial \alpha}\right)^2, \tag{6.98}$$

where the integral runs over the modes extracted, and

$$\frac{1}{\sqrt{I(\alpha)}} =: \Delta(\alpha) \tag{6.99}$$

can be thought of as approximating the constraints on α achievable with these modes. We focus for reasons that become clear below primarily on the parameter $\ln \sigma_8^2$, which has

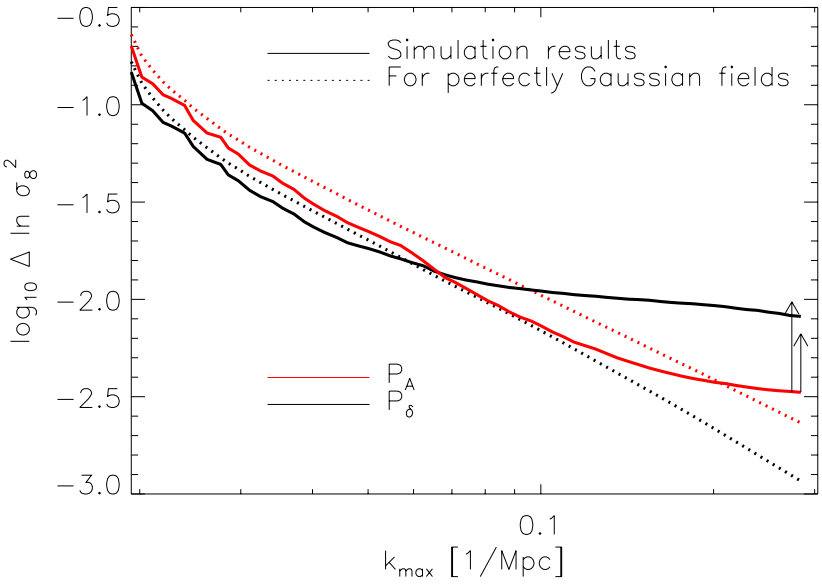

Fig. 6.7 — Comparison of various estimates of the error bar on the linear power spectrum amplitude, $\ln \sigma_8^2$, constrained using power spectra of the overdensity δ (black) and the log-density A (red) in an N-body simulation. Solid curves show how the error bars tighten as the maximum k analyzed increases up to the Nyquist frequency, as in e.g. Neyrinck (2011), equation (6.100). Dotted curves neglect the non-Gaussian component of the covariance matrices, as well as the discrete nature of the Fourier-space mode lattice, equation (6.98). The arrows (one for each choice of σ_A, 0.7 and 0.9) show the expected degradation of the error bars from analyzing δ instead of A in our model given by equation (6.110); these factors appear numerically in the first column of Table 6.1.

a roughly constant impact both on $\ln P_\delta$ and $\ln P_A$. In figure 6.7, we compare this for the δ field and the A field as function of $k_{\rm max}$. The solid lines are the simulation results, evaluating the covariance matrix $C_{kk'}$ between the modes k and k' and setting

$$\Delta(\ln \sigma_8^2) = \sum_{k,k' \leq k_{\rm max}} \frac{\partial \ln P(k)}{\partial \ln \sigma_8^2} C^{-1}_{kk'} \frac{\partial \ln P(k')}{\partial \ln \sigma_8^2}, \tag{6.100}$$

while the dashed lines are in both cases equation (6.99) given by (6.98), with the derivatives being those extracted from the simulations. Since the derivatives are roughly constant, the dashed lines scale like $k^{-3/2}$, i.e. the inverse root of the number of modes. It is clear that the log transform extends the (rough) validity of the Gaussian approximation in terms of Fisher information to the full range of scales we are dealing with. Note however that this is a statement only up to the four point level, since those are the only ones that enter (6.98) and (6.100).

To compare these results to analytical predictions from lognormal statistics, we first

note the following. For a parameter, such as $\ln \sigma_8^2$, that obeys roughly

$$\frac{\partial \ln P_A(k)}{\partial \alpha} \approx \text{cst} =: c, \tag{6.101}$$

the correlated Gaussian field A is equivalent, from the point of the view of the information on that parameter, to a field with the same variance but with $\xi(r) = 0$ for $r > 0$. This may not sound like an obvious statement so let us show this explicitly : start from equation (6.98) which leads to

$$I(\alpha) = c^2 \, \frac{V}{2} \int \frac{d^3k}{(2\pi)^3} \tag{6.102}$$

The integration on the right, in a discrete description, is the number of available modes, equal to the number d of grid points, times the spacing of the modes $\Delta k = (2\pi)^3/V$. It follows

$$I(\alpha) = c^2 \, \frac{d}{2}. \tag{6.103}$$

On the other hand, the observation of d uncorrelated Gaussian variables with variance σ_A^2 always carries the information

$$d \left(\frac{\partial \sigma_A^2}{\partial \alpha} \right)^2 \frac{1}{2\sigma_A^4} \tag{6.104}$$

in their variances. If the derivative of $\ln P$ is the constant c, we have

$$\begin{aligned} \frac{\partial \sigma_A^2}{\partial \alpha} &= \int \frac{d^3k}{(2\pi)^3} \frac{\partial P_A(k)}{\partial \alpha} \\ &= \int \frac{d^3k}{(2\pi)^3} P_A(k) \frac{\partial \ln P_A(k)}{\partial \alpha} \\ &= c \, \sigma_A^2. \end{aligned} \tag{6.105}$$

and thus expressions (6.103) and (6.104) are identical. In terms of information on such parameters, the correlated, Gaussian A field is thus exactly equivalent to d uncorrelated Gaussian variables with the same variances. These parameters can be seen as entering therefore predominantly the variance, the two point correlation function at zero lag, that contains most information, and the correlations at non zero lag carrying little independent information. This is also expected to hold for the δ field, since it is very non-linear and the variance dominates over the clustering in the two point correlation matrix, i.e. the two point correlation matrix is close to diagonal, so that the variance will dominate in any covariance matrix, as well as in the sensitivity to the parameter.

Since information just adds up for any number independent variables, this means that we can try and use directly the exact results we derived in the previous section for the one dimensional lognormal distribution to get a rough but still reasonable estimate of the improvement in the constraints from analyzing the A field. In that section were derived all necessary quantities to obtain the cumulative efficiencies

$$\epsilon_N(\alpha) := \frac{1}{I(\alpha)} \sum_{n=2}^{N} s_n^2(\alpha) \in (0, 1) \tag{6.106}$$

of the first N moments of the δ field to catch the information in A. These coefficients are extremely sensitive functions of σ_A^2, decaying like $\exp(-4\sigma_A^2) \sim \sigma_\delta^{-8}$ as soon as σ_A becomes close to unity.

There is a slight modification to make to these coefficients so that we can confront them to the simulations results. From the simulations only the spectrum of A were extracted, but not the mean of A, which also carries information in principle, even if δ itself has zero mean. For a one dimensional lognormal variable with unit mean, we have from equation (6.9) that $\bar{A} = -\frac{1}{2}\sigma_A^2$. For that lognormal variable the total information is given by the usual formula for the Gaussian A,

$$I(\alpha) = \frac{1}{\sigma_A^2}\left(\frac{\partial \bar{A}}{\partial \alpha}\right)^2 + \frac{1}{2\sigma_A^4}\left(\frac{\partial \sigma_A^2}{\partial \alpha}\right)^2. \tag{6.107}$$

It reduces thus to

$$I(\alpha) = \frac{1}{2\sigma_A^4}\left(\frac{\partial \sigma_A^2}{\partial \alpha}\right)^2 \left(1 + \frac{\sigma_A^2}{2}\right), \tag{6.108}$$

where the rightmost term contains the part of the information in the mean of A. The efficiencies ratios of the moments of δ to that of the variance of A only, excluding the mean, becomes thus

$$\tilde{\epsilon}_N(\alpha) := \epsilon_N(\alpha)\left(1 + \frac{\sigma_A^2}{2}\right). \tag{6.109}$$

Note that in principle theses efficiencies $\tilde{\epsilon}(\alpha)$ can now be larger than unity, if the moments of δ would capture not only the information in σ_A^2, but also that in $\bar{A}$.

The improvement factors, i.e. the ratio of the constraints on α from analyzing the first N correlation functions of δ, to the the constraint from the two-point function of A, are thus in this model

$$[\tilde{\epsilon}_N(\alpha)]^{-1/2} =: \Delta_N^\delta(\alpha)/\Delta_2^A(\alpha). \tag{6.110}$$

They are independent of the parameter α in this one dimensional picture, since the only relevant parameter is σ_A^2, or equivalently σ_δ^2. Remember that the denominator on the right hand side can actually be calculated for any lognormal field from (6.98), our additional assumptions can be seen thus as entering only the numerator. We argued that these ratios are expected to be roughly correct for parameters such as $\ln \sigma_8^2$, but they become in all cases exact for a lognormal field whose variance dominates enough the clustering, $\xi_\delta(r)/\sigma_\delta^2 \ll 1$, for all r. The effective nearest neighbor distance given the modes we used can be evaluated as $r_{\min} \approx \int \frac{d^3k}{(2\pi)^3}^{-1/3}$, and we find $\xi_\delta(r_{\min})/\sigma_\delta^2 = 0.3$.

Finally, there is slight ambiguity in evaluating $\tilde{\epsilon}_n(\alpha)$. A purely lognormal field has $\sigma_A = [\ln(1+\sigma_\delta^2)]^{1/2}$, but this relation is not fulfilled precisely in our simulations. We obtain $\sigma_A = 0.7, \sigma_\delta = 1.1$ and so , $[\ln(1+\sigma_\delta^2)]^{1/2} = 0.9$ rather than 0.7. This discrepancy may be due of course to an intrinsic failure of the lognormal assumption, or to the presence of the smallest scales, slightly correlated, as seen from the start of saturation in figure 6.7.

We show in the first two rows of Table 6.1 the factors of improvement for these two

Tab. 6.1 — The factors of improvements in constraining power of the spectrum of the log-density field in the lognormal model and as seen in simulations. See text for more details.

	$\Delta^\delta_2/\Delta^A_2$	$\Delta^\delta_3/\Delta^A_2$	$\Delta^\delta_\infty/\Delta^A_2$
LN, $\sigma_A = 0.7$	2.0	1.6	1.3
LN, $\sigma_A = 0.9$	2.9	2.4	2.1
Sim. $\alpha = \ln \sigma_8^2$	2.5		
Sim. $\alpha = n_s$	2.4		

values of $\sigma_A, 0.7$ and 0.9, for $N = 2, 3$ and ∞. In the third row is shown the improvement found extracting P_A rather than P_δ in the simulations. Given our assumptions, and the very high sensitivity of $\epsilon_N(\alpha)$ to the variance of the field, they agree remarkably : for the sake of comparison, a variance twice as large of $\sigma_A = 2 \to \sigma_\delta = 7.3$ would have predicted a factor of $\Delta_2(\delta)/\Delta_2(A) = 522$, and for $\sigma_A = 3 \to \sigma_\delta = 90$ a factor of $\approx 5 \cdot 10^6$.

We also performed this analysis for the tilt parameter n_s, which from its very definition has a very differentiated impact over different modes, and finding, just as in the original analysis (Neyrinck, 2011), that the improvement factor is roughly parameter independent as shown in the fourth row of the table. This is another argument supporting the view that the dynamics of the information are indeed captured by such a simple picture. It may be due to the fact that the smallest scales, containing the largest number of modes, contributes the majority of the information in A for any parameter, and thus that the sensitivity can be effectively treated as constant, equal to its value on small scales, making our argument above valid for basically any parameter. Note that for both values of σ_A the spectrum of A still outperforms the entire hierarchy of δ by a sizeable factor for the lognormal model. Of course, this is much more speculative.

As a concluding remark, we did not consider observational noise issues in this section. It remains therefore unclear to what extent these improvements can be achieved with actual galaxy survey data. Generically, it is reasonable to expect that noise will reduce these improvement factors. This section nonetheless makes clear that in this case, improving the specifications of a survey in order to decrease the observational (e.g. shot) noise will be at the same time actually reducing the efficiency with which cosmological parameters can be extracted with the hierarchy of δ (i.e. the fraction of information that is contained in the hierarchy with respect to the total). First elements towards noise issues in this context are presented in section 6.6.2

6.6 Information at all orders for uncorrelated fiducial

In this section, we extend the results of section 6.4, going to the next simplest step. In section 6.4 we derived the information coefficients at all order for the uncorrelated log-

normal field. Now we solve the situation of an uncorrelated fiducial. That is to say we allow parameters to create correlations between the different variables according to the lognormal score functions, while keeping the fiducial model model uncorrelated. In this case, all correlations contribute in general to the information. This is equivalent to use the exact derivatives of the N-point moments, while keeping their covariance matrices as those of the uncorrelated model.

Of course, this is still a very much simplified picture. Nevertheless it corrects one of the defect of the model we used earlier to compare to the N-body simulations, namely that its predictions are model parameter independent, since the only relevant parameter was the variance. In the following differentiated impact on the two-point function allow to make different predictions for different parameters. Also, it provides us for the first time with an easy to implement model for the statistical power of the point functions at all orders, where N-point moments at nonzero lag genuinely contribute to the information. This model can also be extended to deal with the presence of noise, see section 6.6.2.

We work throughout this section first in d dimensions, and then extend these results to the continuous field limit, in Fourier space notation. We assume throughout statistical homogeneity and isotropy. Also, since we will be interested to compare this to the simulations we assume again that the mean of the field is independent of the parameter. In this case, we have $\partial_\alpha \bar{A} = -\partial_\alpha \sigma_A^2/2$.

A parameter α can enter the two-point correlation function at any argument. As in previous sections $p_\rho^{LN}(\rho, \boldsymbol{\theta})$ is the d-dimensional lognormal density function for $\rho = (\rho_1, \cdots, \rho_d)$, and $p_\rho^{LN}(\rho_i, \boldsymbol{\theta})$ are the one dimensional marginals, identical for all i.

Remember that for an uncorrelated fiducial, a most convenient choice of orthogonal polynomials, that we adopt in the following, is given by

$$P_{\mathbf{n}}(\rho) := \prod_{i=1}^{d} P_{n_i}(\rho_i), \tag{6.111}$$

where in our case P_n are the Stieltjes-Wigert polynomials introduced earlier in 6.4. We will therefore be able to express our results in terms of those for the one dimensional distribution in that section.

From the definition of lognormal variables, the score function of the d dimensional distribution reads

$$\frac{\partial \ln p_\rho^{LN}(\rho, \boldsymbol{\theta})}{\partial \alpha} = -\frac{1}{2} \sum_{i,j=1}^{d} \left(A - \bar{A}\right)_i \frac{\partial \left[\xi_A^{-1}\right]_{ij}}{\partial \alpha} \left(A - \bar{A}\right)_j - \frac{1}{2} \frac{\partial \sigma_A^2}{\partial \alpha} \sum_{i,j=1}^{d} \left[\xi_A^{-1}\right]_{ij} (A - \bar{A})_j + \text{cst.} \tag{6.112}$$

Under our assumption of uncorrelated fiducial, we have

$$\frac{\partial \xi_A^{-1}}{\partial \alpha} = -\xi_A^{-1} \frac{\partial \xi_A}{\partial \alpha} \xi_A^{-1} = -\frac{1}{\sigma_A^4} \frac{\partial \xi_A}{\partial \alpha}. \tag{6.113}$$

The easiest way to obtain the information coefficients $s_{\mathbf{n}}(\alpha)$ is to rewrite this last expression in terms of the score function of the one dimensional distribution, for which we have already calculated the coefficients for $\alpha = \ln\bar{\rho}$ and $\alpha = \sigma_\delta^2$.

Separating diagonal and non diagonal elements of ξ_A we obtain after a straightforward calculation

$$\frac{\partial \ln p_\rho^{LN}(\rho,\boldsymbol{\theta})}{\partial\alpha} = \frac{\partial\sigma_\delta^2}{\partial\alpha}\sum_{i=1}^{d}\frac{\partial \ln p_\rho^{LN}(\rho_i,\boldsymbol{\theta})}{\partial\sigma_\delta^2} + \frac{1}{2}\sum_{i\neq j=1}^{d}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\alpha}\frac{\partial \ln p_\rho^{LN}(\rho_i,\boldsymbol{\theta})}{\partial\ln\bar{\rho}}\frac{\partial \ln p_\rho^{LN}(\rho_j,\boldsymbol{\theta})}{\partial\ln\bar{\rho}}. \tag{6.114}$$

We can now evaluate easily the information coefficients. Consider a fixed order $N = |\mathbf{n}|$. From the product form of the polynomials and the fact that the score function couples variables at most in pairs, we can conclude that

$$s_{\mathbf{n}}(\alpha) = 0, \quad \text{if } \mathbf{n} \text{ has more than two nonzero indices.} \tag{6.115}$$

Then, using the fact that $\langle\partial_\alpha \ln p\rangle = 0$ for any parameter (normalisation of the density), it follows that for $\mathbf{n}$ with only one non zero index the second term in (6.114) does not contribute to $s_{\mathbf{n}}(\alpha)$. We obtain

$$s_{\mathbf{n}}(\alpha) = \frac{\partial\sigma_\delta^2}{\partial\alpha}s_N(\sigma_\delta^2), \text{ for } \mathbf{n} = (0,\cdots,0,N,0,\cdots,0), \tag{6.116}$$

wherever the nonzero entry is. Finally, for $\mathbf{n}$ with two nonzero entries we have this time that the first term does not contribute, with the result

$$s_{\mathbf{n}}(\alpha) = \frac{\partial\left[\xi_A\right]_{ij}}{\partial\alpha}s_N(\ln\bar{\rho})s_{N-k}(\ln\bar{\rho}), \text{ for } \mathbf{n} = (0,\cdots,0,k,0,\cdots,0,N-k,0,\cdots,0), \tag{6.117}$$

and i,j are the indices of the nonzero entries. Summing up over all such multiindices of order N, we obtain

$$\left[F_N\right]_{\alpha\beta} = d\,\frac{\partial\sigma_\delta^2}{\partial\alpha}\frac{\partial\sigma_\delta^2}{\partial\beta}s_N^2(\sigma_\delta^2) + \left(\frac{1}{2}\sum_{i\neq j=1}^{d}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\alpha}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\beta}\right)\sum_{i=1}^{N-1}s_N^2(\ln\bar{\rho})s_{N-i}^2(\ln\bar{\rho}). \tag{6.118}$$

6.6.1 Cumulative efficiencies

For our uncorrelated fiducial, we have from (6.13) that the total information is given by

$$\begin{aligned} F_{\alpha\beta} &= d\frac{\partial\sigma_A^2}{\partial\alpha}\frac{\partial\sigma_A^2}{\partial\beta}\left(\frac{1}{4\sigma_A^2}+\frac{1}{2\sigma_A^4}\right) + \frac{1}{2\sigma_A^4}\sum_{i\neq j=1}^{d}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\alpha}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\beta} \\ &:= F_{\alpha\beta}^{\sigma^2} + F_{\alpha\beta}^{\xi}. \end{aligned} \tag{6.119}$$

Summing the information coefficient over N and dividing by the total information content, we obtain the corresponding cumulative extraction efficiencies :

$$\left[F_{\leq N}\right]_{\alpha\beta} = F_{\alpha\beta}^{\sigma^2}\,\epsilon_N(\sigma_\delta^2) + F_{\alpha\beta}^{\xi}\,\epsilon_N(\xi), \tag{6.120}$$

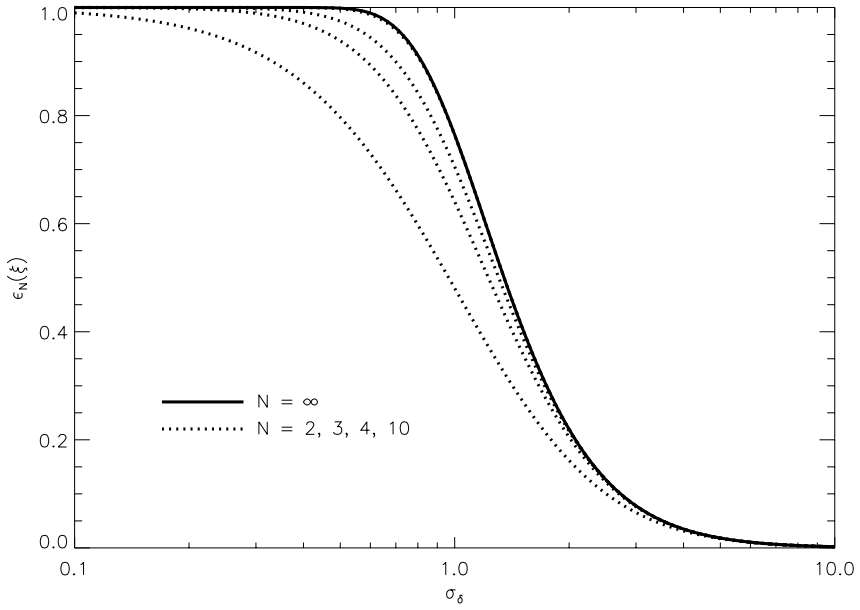

Fig. 6.8 — The fraction of the total information content that is accessible through extraction of the full series of N-point moments of the lognormal field for an uncorrelated fiducial model, for a parameter entering the two-point function at non zero lag, but not the variance.

where $\epsilon_N(\sigma_\delta^2)$ was defined in section 6.4 (and shown in the right panel of figure 6.4), and the second efficiency is

$$\epsilon_N(\xi) := \sigma_A^4 \sum_{n=2}^{N} \sum_{i=1}^{N-1} s_i^2(\ln \bar{\rho}) s_{n-i}^2(\ln \bar{\rho}) \in (0,1). \tag{6.121}$$

Figure 6.8 shows the behavior of $\epsilon_N(\xi)$ as function of σ_δ. It is obviously qualitatively very similar to figure 6.4 for $\epsilon(\sigma_\delta^2)$, showing a slower rate of decay to zero though.

6.6.2 Impact of discreteness effects

It is a very interesting to note that these exact results for an uncorrelated fiducial can be extended with little additional (numerical) effort to the case of a noisy field.

Consider that the observed field $\rho_{\rm obs}$ is not exactly the lognormal density field but a noisy tracer :

$$p_{\rho_{\rm obs},\boldsymbol{\theta}}(\rho_{\rm obs}) = \int d^d\rho \, p_\rho^{LN}(\rho,\boldsymbol{\theta}) \prod_{i=1}^{d} p_{\rm noise}(\rho_{{\rm obs},i}|\rho_i), \tag{6.122}$$

for some one dimensional, parameter independent, density function $p_{\rm noise}(\rho_{\rm obs}|\rho)$. It can represent for instance additive noise, or discreteness effects if $\rho_{\rm obs}$ is for instance a number

of dark matter particles, or galaxies, in a cell of a given volume. For an uncorrelated fiducial, direct calculation shows that equation (6.114) stays in fact formally unchanged :

$$\begin{aligned}\frac{\partial \ln p_{\rho_{\rm obs}}(\rho,\boldsymbol{\theta})}{\partial\alpha} &= \frac{\partial\sigma_\delta^2}{\partial\alpha}\sum_{i=1}^{d}\frac{\partial \ln p_{\rho_{\rm obs}}(\rho_{{\rm obs},i},\boldsymbol{\theta})}{\partial\sigma_\delta^2}\\ &+\frac{1}{2}\sum_{i\neq j=1}^{d}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\alpha}\frac{\partial \ln p_{\rho_{\rm obs}}(\rho_{{\rm obs},i},\boldsymbol{\theta})}{\partial\ln\bar{\rho}}\frac{\partial \ln p_{\rho_{\rm obs}}(\rho_{{\rm obs},j},\boldsymbol{\theta})}{\partial\ln\bar{\rho}}.\end{aligned} \tag{6.123}$$

It follows that the results (6.116) and (6.117) for the information coefficents, as well as the information of order N, equation (6.118), that we derived in the last section hold unchanged, with the understanding that the coefficients s_n present there are not those of the one dimensional lognormal density function but now those of the one dimensional density for $\rho_{\rm obs}$. Those coefficients can generically be obtained numerically without much difficulty even for high orders.

The total information becomes from (6.123)

$$F_{\alpha\beta} = d\,F^{1D}_{\alpha\beta} + \frac{1}{2}\sum_{i\neq j=1}^{d}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\alpha}\frac{\partial\left[\xi_A\right]_{ij}}{\partial\beta}\left[F^{1D}_{\ln\bar{\rho}\ln\bar{\rho}}\right]^2, \tag{6.124}$$

where F^{1D} is the information matrix of the 1 dimensional density for $\rho_{\rm obs}$, that can also be evaluated numerically for a prescribed shape of $p_{\rm noise}$.

We evaluated how the efficiencies change from the purely lognormal case, figure 6.4 and 6.8, when including discreteness effects in the form of the Poisson model, introduced already earlier in 3.2.3. In this case, the observations are Poisson variables, with intensity in each cell given by the value of the lognormal field in that cell. The one-point density for $\rho_{\rm obs} =: n$ object in a cell is given by

$$p_n(n,\boldsymbol{\theta}) = \int d\rho\, p^{LN}_\rho(\rho,\boldsymbol{\theta})e^{-\rho n_c}\frac{(\rho n_c)^n}{n!},\quad n=0,1\cdots. \tag{6.125}$$

Here the lognormal distribution is set to have mean unity, so that n_c is the mean number of particles per cell. We already know from section 3.2.3 that the matrices $F_{\leq N}$ must always be reduced, as well as the total information. Of course, they might be reduced at a different rate, so that the efficiencies need not be reduced, but can increase drastically, if the observed density function is made closer to a Gaussian.

The results are shown on figure 6.10. It is shown, as function of n_c, the total information F, solid line, as well as $F_{\leq N}$, dashed lines, all normalised by the total information content of the noise free lognormal distribution, at a fixed variance $\sigma_\delta^2 = 0.7$. The upper panel is for a parameter that enter the variance only and not the correlations, and the lower panel for a parameter that enter the correlations exclusively. It is clear that the strongest effect is to reduce the total information content, as soon as the mean number of objects is close to unity. It follows that in the very noisy regime, the N-point moments are efficient again, in that the fraction of the information that they capture is close to unity

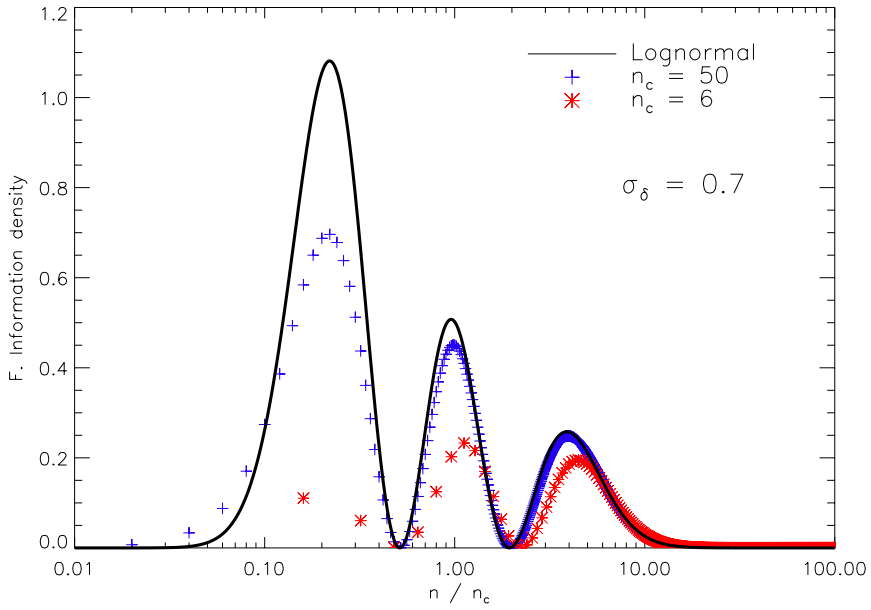

Fig. 6.9 — The Fisher information density of the Poisson sample of the lognormal distribution with $\sigma_\delta = 0.7$, as given in (6.125), for two different values of n_c as indicated. The solid line is that of the (continuous) lognormal distribution. The vast majority of the information was initially within the underdense regions but is destroyed if the mean number of object per cell is low, making the N-point moments efficient again.

again.

We have already argued that the reason for which the N-point moments of the lognormal field are inefficient is that they are unable to probe the underdense regions, rich in information. This is seen very clearly in figure 6.9. Shown are the Fisher information density for the parameter σ_δ^2 of the one dimensional distribution $p_n(n)$, for a mean number of objects n_c of 50 (crosses) and 6 (diamonds). The solid line shows that of the noise-free, continuous lognormal. Unsurprisingly, it is the information in the underdense regions ($n/n_c < 1$), precisely the information that the moments were unable to catch, that is mostly destroyed by discreteness effects. On the other hand, the information in the overdense regions ($n/n_c > 1$), accessible to the moment hierarchy, is left untouched.

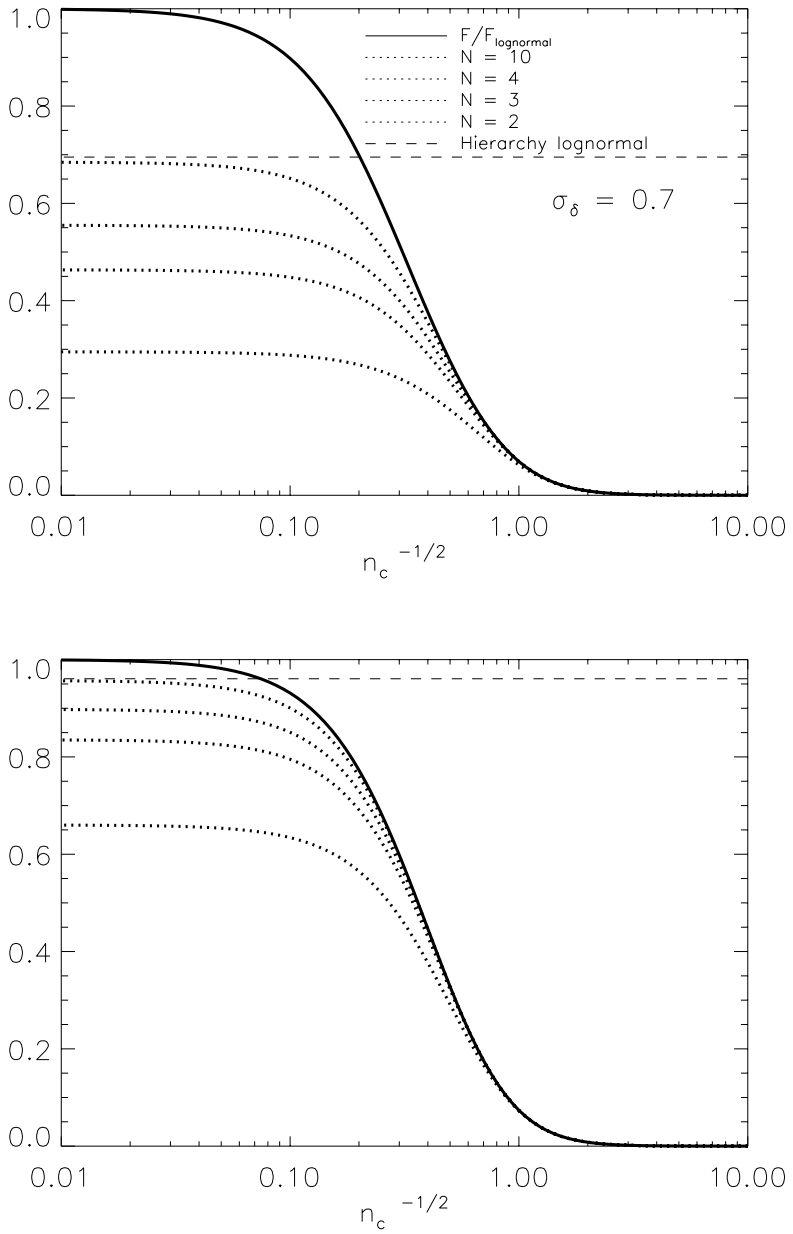

Fig. 6.10 — Upper panel : the information content of the first N-point moments in Poisson sample of a lognormal field with uncorrelated fiducial with $\sigma_\delta = 0.7$, as function of the inverse root of the mean number of objects per cell, for a parameter entering the variance of the field but not the correlations. The solid line denotes the total information content, and the dotted lines that of the point functions up to order N. All lines are normalised to the total information content of the noise-free lognormal field. Dashed is that of the entire hierarchy of the lognormal for comparison. Lower panel : same for a parameter entering correlations but not the variance. N-point functions become efficient again in the noisy regime, but only due to the massive decrease in the total information content.

6.6.3 The reason for a roughly parameter independent factor of improvement

We now go back to our exact expressions (6.120) for the information in N-point moments of a lognormal field with uncorrelated fiducial. We seek to transform these expressions in Fourier space such that it is easier to compare or use N-body simulation results within this model.

Within our assumptions, the fiducial power spectrum is constant and directly proportional to the variance, such that we can write

$$\frac{\partial \ln \sigma_A^2}{\partial \alpha} = \frac{1}{d} \sum_{\mathbf{k}} \frac{\partial \ln P_A}{\partial \alpha}. \tag{6.126}$$

Turning the sum over the d available modes, with spacing $\Delta k = (2\pi)^3/V$ into an integral, it becomes

$$\frac{\partial \ln \sigma_A^2}{\partial \alpha} = V_{\rm cell} \int_0^{k_{\rm max}} \frac{dk\, k^2}{2\pi^2} \frac{\partial \ln P_A}{\partial \alpha}. \tag{6.127}$$

In this equation, $V_{\rm cell}$ is the volume associated to a real space cell, given formally by

$$V_{\rm cell} = \frac{(2\pi)^3}{d\Delta k} = \frac{V}{d} \tag{6.128}$$

that can be written in the continuous notation as

$$\frac{1}{V_{\rm cell}} = \int_0^{k_{\rm max}} \frac{dk\, k^2}{2\pi^2}. \tag{6.129}$$

Thus, we can write the first term in (6.120) as

$$\frac{d}{2} \frac{\partial \ln \sigma_A^2}{\partial \alpha} \frac{\partial \ln \sigma_A^2}{\partial \beta} = V_{\rm cell} \frac{V}{2} \int_0^{k_{\rm max}} \frac{dk\, k^2}{2\pi^2} \frac{\partial \ln P_A(k)}{\partial \alpha} \int_0^{k_{\rm max}} \frac{dk\, k^2}{2\pi^2} \frac{\partial \ln P_A(k)}{\partial \beta}. \tag{6.130}$$

On the other hand, to write the second term of equation (6.120) as an integral over the power spectrum, we first complete the sum over $i \neq j$, and remember that within our assumptions, ξ_A^{-1} is $1/\sigma_A^2$ times the identity matrix. We obtain

$$\frac{1}{2\sigma_A^4} \sum_{i \neq j} \frac{\partial \xi_{ij}}{\partial \alpha} \frac{\partial \xi_{ij}}{\partial \beta} = \frac{1}{2} {\rm Tr} \left[\frac{\partial \xi_A}{\partial \alpha} \xi_A^{-1} \frac{\partial \xi_A}{\partial \beta} \xi_A^{-1} \right] - \frac{d}{2} \frac{\partial \ln \sigma_A^2}{\partial \alpha} \frac{\partial \ln \sigma_A^2}{\partial \beta}. \tag{6.131}$$

The first term on the right is nothing else than the well known expression for the information content of the two point correlation function of the Gaussian field A. Fourier transformation of the first term leads to

$$\frac{V}{2} \int_0^{k_{\rm max}} \frac{dk\, k^2}{2\pi^2} \frac{\partial \ln P_A(k)}{\partial \alpha} \frac{\partial \ln P_A(k)}{\partial \beta} =: F_{\alpha\beta}^{P_A}. \tag{6.132}$$

This last expression (6.132), the information content of the spectrum of the Gaussian field A, is exact for any lognormal field, the assumptions we are making above not entering it.

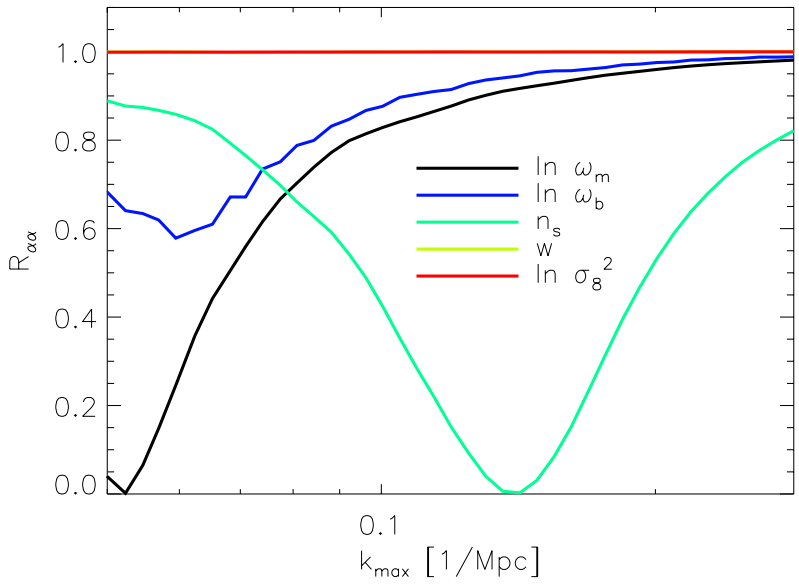

Fig. 6.11 — The ratios $R_{\alpha\alpha}$, equation (6.133), for the five cosmological parameters studied in the analysis of Neyrinck (2011), as function of the smallest scale present in the analysis. The constancy of these factors once all modes are included implies a generic, roughly parameter independent gain from the analysis of the spectrum of A rather than δ. See text for more details.

We are now ready to put all these relations together. In section 6.5 (see also Neyrinck (2011)) we just studied the constraints on the parameters given by extracting the spectrum of A or δ. We already know what the information in the spectrum of A is, equation (6.132). We will therefore find convenient to consider the ratios $F_{\leq N}/F^{P_A}$. Define first the ratios $R_{\alpha\beta}$ as

$$R_{\alpha\beta} := \frac{V_{\rm cell}}{2\pi^2} \frac{\int_0^{k_{\rm max}} dk\, k^2 \partial_\alpha \ln P_A \int_0^{k_{\rm max}} dk\, k^2 \partial_\beta \ln P_A}{\int_0^{k_{\rm max}} dk\, k^2 \partial_\alpha \ln P_A \partial_\beta \ln P_A}. \tag{6.133}$$

These ratios are unity for parameters that obeys $\partial_\alpha \ln P_A \approx$ cst, such as $\ln \sigma_8^2$. They can vanish for parameters with differentiated impact, such as the spectral index n_s.

We obtain from (6.120) and the relations in this section that the loss of information by extracting the first N correlation functions of δ rather than the spectrum of A is given by

$$\frac{[F_{\leq N}]_{\alpha\beta}}{F^{P_A}_{\alpha\beta}} = \epsilon_N(\xi)\,[1 - R_{\alpha\beta}] + \epsilon_N(\sigma_\delta^2) R_{\alpha\beta} \left(1 + \frac{1}{2}\sigma_A^2\right). \tag{6.134}$$

Recall that while the denominator in (6.134) equation (6.132) is exact for any lognormal field, the numerator must be exact only within the assumption of an uncorrelated fiducial

as we specified. All the terms of this equation are strong functions of the variance of the field, or equivalently of $k_{\max}$. The variance of A being set by

$$\sigma_A^2(k_{\max}) = \int_0^{k_{\max}} \frac{dk\, k^2}{2\pi^2} P_A(k). \tag{6.135}$$

We are now in position to test whether our more sophisticated model changes our conclusions from section 6.5, where we reproduced the correct factor of improvements with our one dimensional model. Of course, it is apparent that the key element is the magnitude of the coefficients $R_{\alpha\alpha}$ in (6.134): if they are close to unity then the prediction of the model is essentially unchanged from that of the one dimensional distribution. These factors, obtained from the derivatives of the spectrum measured from the simulations as in section 6.5, are shown in figure 6.11 as function of $k_{\max}$ for the parameters $\omega_b = h^2\Omega_b$, $\omega_m = h^2\Omega_m$, the tilt n_s, as well as the dark energy equation of state w and $\ln \omega_8^2$. The last two curves are indistinguishable from unity on this figure.

It is obvious that when all the scales are included, these coefficients are all very close unity, leading to nearly identical predictions for the improvement factors for all these parameters. The tilt parameter is the one that differs the most, where the prediction in this improved model is slightly below the one dimensional case, consistently with the findings in simulations that we presented in section 6.5(see table 6.1). The other parameters show essentially identical dynamics, and the predictions do not differ from the one dimensional case. This is also remarkably consistent with the study of Neyrinck (2011), where n_s is the only parameter to show a slightly different behavior (remember though that that work includes some modes $0.3/\text{Mpc} < k < 0.6/\text{Mpc}$ that cannot be described as lognormal).

Maybe surprisingly, our simple, essentially one dimensional treatment thus definitely seems to capture very well the dynamics of the information as seen in N-body simulations.

6.7 The general case : a difficult Lagrange interpolation problem

We end this chapter by a remark on the problem of the derivation the matrices F_N in the case of a lognormal field with arbitrary correlations, a still unsolved problem, except in the situations dealt with in section 6.4 and 6.6.

Remember that we could solve the one dimensional case making use of a curious identity of the polynomials orthogonal P_n to the lognormal distribution. We showed that the polynomial in t defined as

$$\langle P_n(t\rho) \rangle \tag{6.136}$$

has zeroes in $q^{-i}, i = 0, n-1$. Since a polynomial in one variable factorizes in product of roots, we could easily obtain explicit form for the polynomials as well as of the information

coefficients.

It is very interesting that this property of the set of orthogonal polynomials generalises to any dimensioniality d. In the multiindex notation, recall that the N-point moments are given by

$$m_{\mathbf{n}} = \langle \rho(x_1)^{n_1} \cdots \rho(x_d)^{n_d} \rangle = \exp\left(\bar{A} \cdot \mathbf{n} + \frac{1}{2} \mathbf{n} \cdot \xi_A \mathbf{n} \right). \tag{6.137}$$

The following property is easy to verify :

$$m_{\mathbf{n}+\mathbf{m}} = m_{\mathbf{n}} m_{\mathbf{m}} Q_{\mathbf{nm}}, \tag{6.138}$$

where we defined

$$Q_{\mathbf{nm}} := \exp\left(\mathbf{n} \cdot \xi_A \mathbf{m} \right). \tag{6.139}$$

This identity is the generalisation of the property $m_{i+j} = m_i m_j q^{-ij}$ which we made good use of in section 6.4. We can write the matrix in an other useful form. Define the points $t_{\mathbf{n}}$ as

$$t_{\mathbf{n}} = \exp\left(\xi_A \cdot \mathbf{n} \right), \text{ i.e. } (t_{\mathbf{n}})_i = \exp\left(\sum_{j=1}^{d} [\xi_A]_{ij} \, n_j \right). \tag{6.140}$$

We have the following relation

$$Q_{\mathbf{nm}} = t_{\mathbf{n}}^{\mathbf{m}}. \tag{6.141}$$

The matrix Q is thus a Vandermonde matrix in several variables.

We can now prove the following property of the orthogonal polynomials in any dimensionality. For any $\mathbf{n}$ and $\mathbf{m}$ holds:

$$\langle P_{\mathbf{n}}(t_{\mathbf{m}} \rho) \rangle = \frac{1}{m_{\mathbf{m}}} \langle \rho^{\mathbf{m}} P_{\mathbf{n}}(\rho) \rangle. \tag{6.142}$$

To prove that relation expand $P_{\mathbf{n}}(\rho) = \sum_{\mathbf{k}} C_{\mathbf{nk}} \rho^{\mathbf{k}}$ on both sides and use (6.138). Remember that the orthogonal polynomials of a given order are defined such that they are orthogonal to all polynomials of lowest order. If for convenience a further ordering of the multindices of same order $|\mathbf{n}|$ is chosen, we can say without ambiguity whether for two multindices $\mathbf{n}$ and $\mathbf{m}$ hold $\mathbf{n} < \mathbf{m}$, $\mathbf{n} > \mathbf{m}$ or $\mathbf{n} = \mathbf{m}$. We can then write

$$\langle P_{\mathbf{n}}(t_{\mathbf{m}} \rho) \rangle = 0, \text{ for } \mathbf{m} < \mathbf{n}. \tag{6.143}$$

This identity is is the direct analog of the relation we used to solve the one dimensional problem as just discussed at the very beginning of this section. We know enough zeroes of the polynomial

$$\pi_{\mathbf{n}}(t) = \langle P_{\mathbf{n}}(t\rho) \rangle \tag{6.144}$$

in $t = (t_1, \cdots, t_d)$, to specify it uniquely, and with it the orthogonal polynomials.

The problem becomes therefore that to find the polynomials in d variables with the given zeroes as prescribed above, which is nothing else than a familiar Lagrange interpolation problem, with associated Vandermonde matrix Q. Unfortunately, this is of course a considerably more involved task in several dimensions than in one, where many properties such as factorization of the polynomials does not hold anymore.

7 Information escaping the correlation hierarchy of the convergence field

The text of this chapter follows that of the letter we published in Carron (2012b). There we evaluated with numerical methods the information content of the moment hierarchy of the noise-free weak lensing convergence field. The approach presented in chapter 3 is applied this time not to the somewhat ad-hoc model of the lognormal distribution of chapter 4, but to the most realistic case of fits to N-body simulations outputs already present in the literature. We find the dynamics of the information within the hierarchy to be qualitatively very similar than for the lognormal, the moments becoming quickly (even quicker than the lognormal) dramatically inefficient. For these reasons, even though this chapter is comparatively short, we believe it to be an important part of this thesis, as it makes clear that the dynamics explored in chapter 4 with the help of the lognormal field generically affect cosmological non linear fields. It if also found that a simple logarithmic mapping makes the moment hierarchy well suited again for parameter extraction.

7.1 Introduction

N-point correlation functions, first introduced in cosmology by Peebles and collaborators to describe the large scale distribution of galaxies (Peebles, 1980), are now ubiquitous in this field. They are at the heart of many cosmological probes like the CMB, galaxy clustering, or notably weak lensing, which was recognized as one of the most promising probe of the dark components of the universe (Albrecht et al., 2006; Bartelmann and Schneider, 2001; Bernstein, 2009; Hu and Jain, 2004), and which traces the cosmological convergence field.

On large scales, or in the linear regime, correlations are a particularly convenient ap-

proach to tackle the difficult problem of statistical inference on cosmological parameters. Indeed, primordial cosmological fluctuation fields are believed to obey Gaussian statistics, and the first two members of the hierarchy, the mean and the two-point correlation function, provide a complete description of such fields. However, much less is known about the pertinence of the correlation hierarchy in the non-linear regime, or on small scales, where in principle a lot of information is contained, if only due to the large number of modes available for the analysis. More elaborated statistical models must be made in this regime. For instance, the statistics of the matter field and its weighted projection the convergence field were shown to be closer to lognormal, at least in low dimensional settings (Coles and Jones, 1991; Das and Ostriker, 2006; Taruya et al., 2002; Taylor and Watts, 2000), though with sizeeable deviations still.

Two effects relevant for statistical inference can in principle play a role entering the non linear regime, departing from Gaussian initial conditions. First, information may propagate to higher order correlators. Second, the correlation function hierarchy may not provide a complete description of the field anymore , so that information escapes the hierarchy. Even though this second possibility was pointed out qualitatively in an astrophysical context already in Coles and Jones (1991), it seems it was not given further attention in the literature. In this *Letter* we show, using accurate fits of the convergence one-point probability density function to numerical simulations (Das and Ostriker, 2006) that the second effect very quickly completely dominates the convergence field, and thus that the hierarchy is not well suited for inference on cosmological parameters anymore.

Fisher information and orthogonal polynomials. The approach is based on decomposing the Fisher's matrix valued information measure in components unambiguously associated to the independent information content of the correlations of a given order. It was recently proposed in Carron (2011), building upon Jarrett (1984). Exact results at all orders were obtained only for the moment hierarchy of a idealized, perfectly lognormal one dimensional variable, where analytical methods could be applied. In cosmology, the Fisher information matrix is widely used for many years now to estimate the accuracy with which cosmological parameters will be extracted from future experiments aimed at some observables (Bernstein, 2009; Hu and Jain, 2004; Tegmark, 1997, e.g.), assuming Gaussian statistics.

For a general probability density function $p(x, \alpha, \beta)$, α, $\beta, \cdots$ any model parameters, its definition is

$$F_{\alpha\beta} = \left\langle \frac{\partial \ln p}{\partial \alpha} \frac{\partial \ln p}{\partial \beta} \right\rangle . \tag{7.1}$$

Its inverse can be seen through the Cramér-Rao bound (Tegmark et al., 1997) to be the best covariance matrix of the relevant parameters achievable with the help of unbiased estimators. The general procedure to decompose the Fisher information content into uncorrelated pieces, corresponding to an orthogonal system, was presented in a statistical journal in Jarrett (1984). When the observables of interest are products of the variables, i.e. moments or more generally correlation functions, the orthogonal system are orthogonal polynomials. It is discussed in detail in an cosmological context in Carron (2011). In

particular, the variables for which the Fisher information content on α is entirely within the first N pieces, such as the Gaussian variables for $N = 2$, are those for which the function $\partial_\alpha \ln p$ entering (7.1), called the score function, is a polynomial of order N in x. In the case of a single variable, the uncorrelated contribution of order N to the Fisher information matrix $F_{\alpha\beta}$ is given by

$$s_N(\alpha)s_N(\beta), \tag{7.2}$$

where the Fisher information coefficients s_N are the components of the score function with respect to the orthonormal polynomial of order N,

$$s_N(\alpha) = \left\langle \frac{\partial \ln p}{\partial \alpha} P_N(x) \right\rangle, \tag{7.3}$$

$$\langle P_n(x)P_m(x)\rangle = \delta_{mn}, \quad n, m \geq 0 \tag{7.4}$$

For any N, the following relation holds

$$\sum_{n=1}^{N} s_n(\alpha)s_n(\beta) = \sum_{i,j=1}^{N} \frac{\partial m_i}{\partial \alpha} \left[\Sigma^{-1}\right]_{ij} \frac{\partial m_j}{\partial \beta}, \tag{7.5}$$

where $m_i = \langle x^i \rangle$ and $\Sigma_{ij} = m_{i+j} - m_i m_j$ is the covariance matrix. The right hand side being the expression describing the Fisher information content of the moments m_1 to m_N. Whether one recovers the full matrix $F_{\alpha\beta}$ with $N \to \infty$ or only parts of it depends on the distribution under consideration. A sufficient condition is that the polynomials P_n form a complete basis set, which is then essentially equivalent to the condition that the distribution can uniquely be recovered from its moments hierarchy (Carron, 2011; Coles and Jones, 1991, and references therein). This and other sufficient criteria for completeness are tightly linked to the decay rate of the probability density function at infinity.

We define the cumulative efficiency ϵ_N of the moments up to order N to capture Fisher information on α as

$$\epsilon_N(\alpha) := \frac{\sum_{n=1}^{N} s_n^2(\alpha)}{F_{\alpha\alpha}}. \tag{7.6}$$

From the Cramér-Rao bound, $\sqrt{\epsilon_N}$ is the ratio of the the best constraints achievable on α with any unbiased estimator to the expected constraints on α from the extraction of the first N moments .

7.2 Fisher information coefficients

We use the fits to simulations from Das and Ostriker (2006), valid down to the arcsecond scales. Initially built to correct for the failure of the lognormal distribution to reproduce the high and low density tails of the convergence κ on a single lens plane, it reproduces

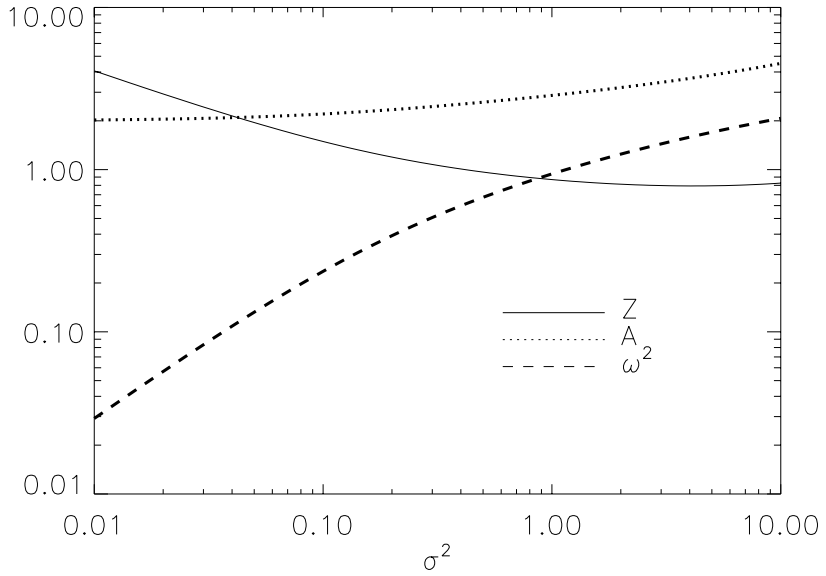

Fig. 7.1 — The three parameters Z,A and ω^2 entering the generalized lognormal model, as function of the variance of $\delta_m^{\rm eff}$.

accurately the cosmological convergence as well, taking into account the broader lensing kernel (Takahashi et al., 2011). In terms of the reduced variable x,

$$x = 1 + \frac{\kappa}{|\kappa_{\rm empty}|} =: 1 + \delta_m^{\rm eff}, \tag{7.7}$$

where $\kappa_{\rm empty}$ is the minimal value of the convergence, corresponding to a light ray traveling an empty region, it takes the form of a generalized lognormal model for the associated effective matter fluctuations $\delta_m^{\rm eff}$,

$$p(x,\sigma) = \frac{Z}{x} \exp\left[-\frac{1}{2\omega^2}\left(\ln x + \frac{\omega^2}{2}\right)^2 \left(1 + \frac{A}{x}\right)\right]. \tag{7.8}$$

In this equation, the three parameters Z, A and ω are such that the mean of x is unity, and its variance $\sigma^2 = \sigma_\kappa^2/\kappa_{\rm empty}^2$ (we are neglecting here a small but non-zero mean of the convergence argued in Takahashi et al. (2011)). Therefore, the only relevant parameter is the variance of the associated matter fluctuations σ^2, fixed by the cosmology from $\kappa_{\rm empty}$ and the convergence power spectrum, together with some filter function corresponding to the smoothing scale, determining the level of non linearity of the field (Das and Ostriker, 2006, figure 1). Linear and non-linear regime being separated at $\sigma^2 \approx 1$. We obtained Z, A and ω^2, shown in figure 7.1, with the help of a standard implementation of the Newton-Raphson method for non-linear systems of equations.

Orthogonal polynomials can very conveniently be generated by recursion, as exposed in details in Gautschi (2004), since they satisfy a three terms recurrence formula. We define

for convenience

$$\hat{\pi}_N(x) := \sqrt{p(x,\sigma)}\pi_N(x), \tag{7.9}$$

where $\pi_N(x)$ is $P_N(x)$ rescaled such that the coefficient of x^N is unity. The recursion relations in Gautschi (2004) become

$$\begin{aligned} \hat{\pi}_{k+1} &= (x-\alpha_k)\hat{\pi}_k - \beta_k\hat{\pi}_{k-1}, \\ \alpha_k &:= \frac{\int_0^\infty dx\, x\, \hat{\pi}_k^2(x)}{\int_0^\infty dx\, \hat{\pi}_k^2(x)} \\ \beta_k &:= \frac{\int_0^\infty dx\, \hat{\pi}_k^2(x)}{\int_0^\infty dx\, \hat{\pi}_{k-1}^2(x)}, \end{aligned} \tag{7.10}$$

and $\pi_{-1}(x) = 0, \pi_1(x) = 1, \beta_0 = 1$, that we implemented using an appropriate discretization of the x-axis. Proper normalization of the polynomials can be performed afterwards. The Fisher information coefficients were then obtained with the help of equation (7.3), using a precise five point finite difference method for the derivatives of Z, A and ω^2 with respect to σ^2 that are needed to obtain the score function.

In figure 7.2, we show the cumulative efficiency $\epsilon_N(\sigma^2)$, for $N = 2$ to $N = 5$, from bottom to top. (Note that $s_1(\sigma^2)$ vanishes since the mean of x is unity for any value of the variance). The uppermost line contains therefore the variance, the skewness, the kurtosis as well as the 5th moment of the field. The contribution of each successive moment can be read out from the difference between the corresponding successive curves. For higher N quick convergence of ϵ_N occurs, presented in figure 7.3 as the solid line, showing ϵ_{10}. For small values of the variance, the field is still close to Gaussian, so that the Fisher information is close to be entirely within the the 2nd moment, and accordingly the ratio ϵ is close to unity in this regime. It is obvious from these figures that the main effect for larger values of the variance is not that Fisher information is transferred to higher order moments, but rather the dramatic cutoff as soon as the variance crosses 0.1. At redshift 1, this corresponds to the scale of $\approx 1'$ (Das and Ostriker, 2006, figure 1), so still within scales probed by weak lensing. For $\sigma \sim 1$, the ratio is close to 0.05, meaning that all moments completely fails to capture the information. Optimal constraints on any cosmological parameter entering σ are thus for this value of the variance a factor $1/\sqrt{0.05} \sim 4.5$ tighter than those achievable with the help of the entire hierarchy.

In figure 7.3 we compare these results to the exact analytical expressions given in Carron (2011) for the lognormal distribution, shown as the dashed line. These are given by, accounting for the different normalization,

$$s_N^2(\sigma^2) = q^2 \frac{q^N}{1-q^N}\left(\prod_{n=1}^{N-1}(1-q^n)\right)\left(\sum_{n=1}^{N-1}\frac{q^n}{1-q^n}\right)^2, \tag{7.11}$$

with $q := 1/(1+\sigma^2)$. The total Fisher information content being in this case $(q/\ln q)^2/2 - q^2/(4\ln q)$. There also the information content of the moments saturates quickly as N grows. It is striking that the incompleteness of the moment hierarchy occurs much earlier in the convergence field than in the lognormal. This can be understood from the following considerations. The main effect of the improved model (7.8) for the convergence is

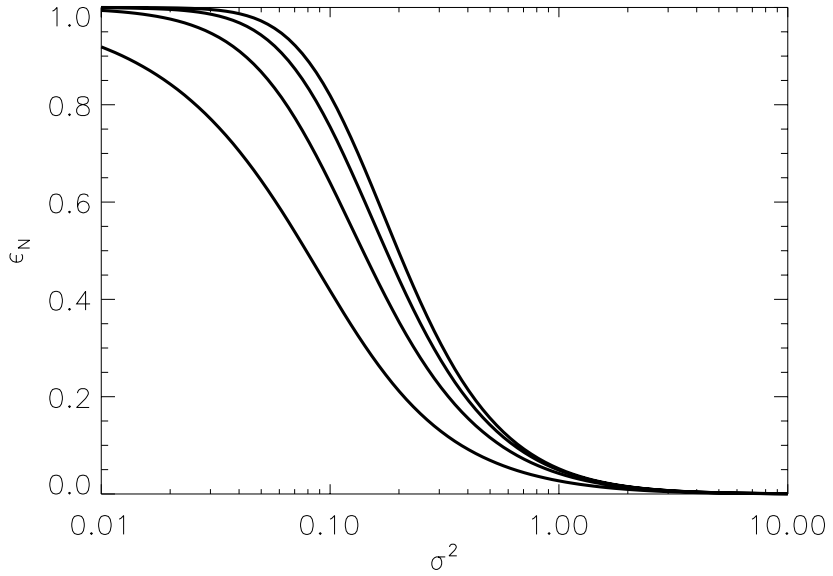

Fig. 7.2 — The cumulative efficiencies ϵ_N of the moments of the convergence in capturing Fisher information, for $N = 2$ to $N = 5$, defined in eq. (7.6), from bottom to top, as function of the variance of $\delta_m^{\rm eff}$.

to reproduce accurately the very sharp cutoff of the probability density function at low convergence values (Das and Ostriker, 2006, figure 3-6). This cutoff is very sensitive to the variance of the field, more sensitive than the cutoff of the lognormal. However, there the contribution to the moment m_n, x^n, is beaten down by orders of magnitude. To make this point clearer, we show in figure 7.4 the Fisher information density $p\,(\partial_{\sigma^2} \ln p)^2$ for the lognormal distribution (dashed) and the model we used (solid), at the scale of $\sigma = 1$ It is obvious in both cases that a large fraction of the information is contained in the underdense regions, describing the cutoff of the distribution, but unaccessible to the moments of x. Since this is even more the case for the convergence field, the efficiency is accordingly even worse.

7.3 Restoration of the information

Finally we investigate to what extent the moment hierarchy of $\ln x$ contains more Fisher information than the hierarchy of x. Though our method is completely independent, this can be seen as complementary to recent works looking at the statistics of the field after local transforms, and at the statistical power of its power spectrum initiated in Neyrinck et al. (2009, 2011); Seo et al. (2011), even though in these works the fact that information actually completely escapes the hierarchy is not appreciated. This is done with the very same method used above, by obtaining the polynomials orthogonal to the distribution of

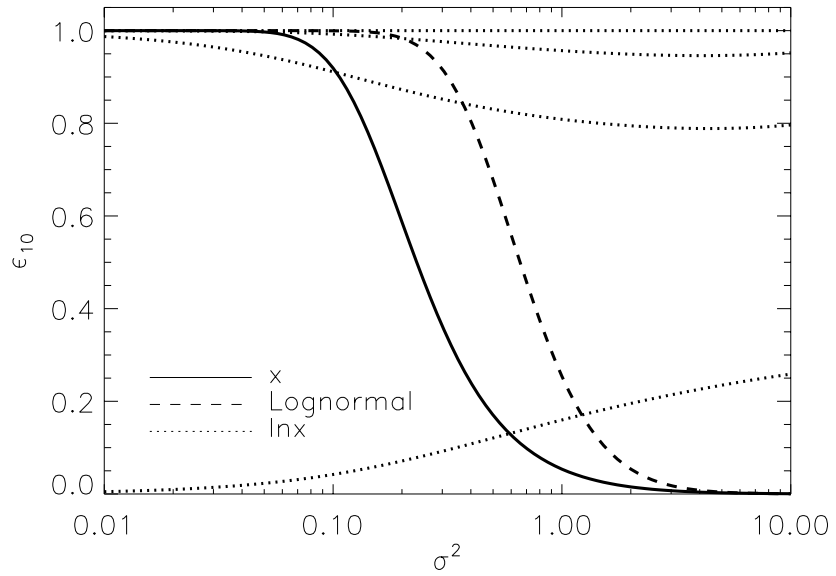

Fig. 7.3 — Solid and dashed line : the efficiency $\epsilon_{N=10}$ of the first 10 moments to capture Fisher information, for the convergence field (solid) and lognormal field (dashed). The curves do not change anymore with increasing N. Dotted : ϵ_1,ϵ_2, ϵ_3 and ϵ_{10} for the logarithmic transform of the field, from bottom to top.

$\ln x$, or equivalently decomposing the score function of x in polynomials in $\ln x$ rather than in x. This is seen to perform very well, as shown by the dotted lines in figure 7.3. From bottom to top are plotted ϵ_1,ϵ_2 and ϵ_3. Also shown in the figure is ϵ_{10} but it is not to be distinguished from unity, meaning that completeness of the hierarchy is restored. We see that over the full range at least 80% of the information is back in the two first moments, and 95% in the first three.

7.4 Conclusions

We have studied the statistical power of the moment hierarchy of the convergence field, when leaving the linear regime. Notably, the hierarchy ceases to provide a complete description of the statistics of the convergence, letting an increasingly large fraction of the Fisher information actually escape the hierarchy, and thus making constraints on cosmological parameters achievable with measurements of the hierarchy suboptimal by increasingly large factors. While our results are exact only for the one point distribution (or equivalently the full correlation function hierarchy of the convergence field in the limit of vanishing correlations), the correlation function hierarchy will also show a similar behavior, though the amplitude of the loss in information and constraining power may vary from parameter to parameter in the details. This is because this defect, for any

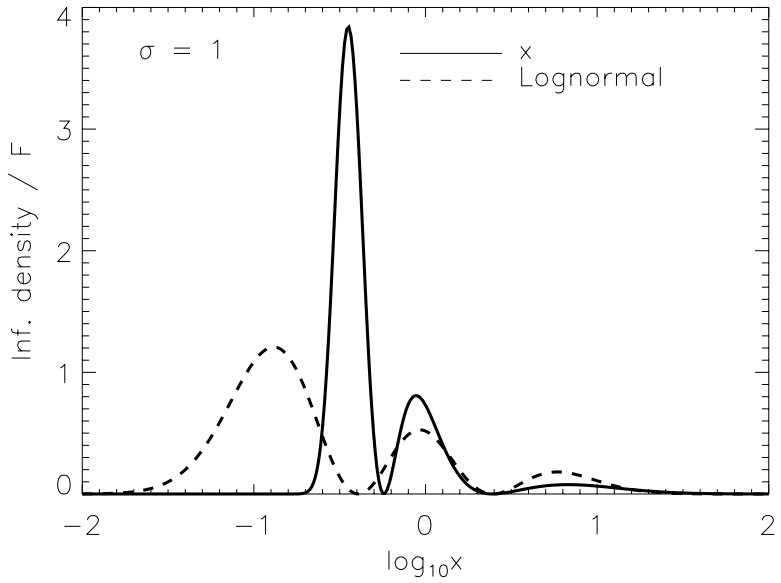

Fig. 7.4 — The Fisher information density of the lognormal (dashed) and the convergence field (solid), renormalized such that it integrates to unity. Clearly, the Fisher information is mostly contained within the underdense regions. The moments are however sensitive to the tail.

number of variables, is due to the very slow decay rate at infinity of the field distribution, which cannot be reproduced by the exponential of a polynomial in the relevant variables. Our findings are consistent with previous analytical results on the lognormal distribution (Carron, 2011), and numerical work from N-body simulations at the power spectrum level (Neyrinck et al., 2009, 2011). Making a tighter connection to such simulation results with the methods presented here is the subject of future work. Of course, the quest for the information in the non linear regime already has problems of its own, such as shot noise issues, or accurate modeling, that we did not consider here. Nonetheless, these results clearly shows that if the correlation function hierarchy is to play a substantial role in getting constraints out of the mildly or non-linear regime, then an approach similar to a Gaussianizing transform (Neyrinck et al., 2011; Seo et al., 2011), in this work a simple logarithmic mapping, can hardly be avoided though the details still needs to be figured out. It is reassuring that this approach seems to work well to first order, and that first steps have recently already been taken in that direction in perturbation theory (Wang et al., 2011), for the matter field. Our work also points toward low convergence regions as carrying large amounts of information, though the importance of noise issues needs to be clarified in this regime. Thus, many promising ways have still to be explored to make profit of mildly and non-linear scales.

Summary of main results and outlook

We have researched the information content of the cosmological matter density field and of the convergence field in order to gain insights on the statistical power of important cosmological observables and their combination, using tools borrowed from information theory.

In a first step, we have introduced distributions of maximum entropy, and discussed with their help the combination of magnification and flexion fields to the more usual shear fields, all tracers of the convergence, for the purpose of extracting cosmological parameters, or to reconstruct the underlying mass field. We found that size information contribute a modest but scale independent amount to the information, and one can expect constraints on any model parameter to increase by some 10% with respect to a shear only analysis. The flexion field has very different noise properties, that are strongly scale dependent. We found the information from flexion alone to take over that of shear only on the scales of the arcsecond, becoming the most interesting observable for the purpose of mass reconstruction. We believe that these scales are rather small in order for flexion to become a useful cosmological tool due to the fact that the nonlinear matter power spectrum is not extremely well understood on these scales at present, but flexion can complement well the shears on the scale of the arcminute.

After a brief discussion of the information content of power spectra estimators in chapter 5, we have then turned to the main part of thesis, which is the study of the statistical power of the hierarchy of N-point moments of a given density function, with applications to the noise-free lognormal and convergence field. We found that the hierarchy is clearly inadequate to describe such fields in the nonlinear regime, as the hierarchy fails to capture increasingly large fractions of the information. The reason is that the information that is captured becomes completely correlated with that of the overdense regions, and the information in the underdense regions, increasingly important, becomes inaccessible. A few large, cosmic variance dominated density peaks dominate the correlations and hide most

of the information, making parameter inference from the N-point moments inefficient by orders of magnitude in the deeply nonlinear regime. The fact that information escapes the hierarchy could also been seen directly from the different fields we defined with the very same hierarchy of N-point moments than the lognormal field. We could successfully confront these results to numerical simulations, and show how simple non linear mappings, already present in the literature, are able to correct for these defects.

An aspect of our results that we found most intriguing is the clear success of the model with uncorrelated lognormal fiducial to reproduce correctly the behavior of the Fisher information that is observed in the N-body simulations for a wide class of model parameters. For all the standard cosmological parameters that we tested in chapter 6, we found that the correlations played very little role in understanding the dynamics of their information. The only relevant parameter was found to be the variance of the field, and not the details of the correlation structure. This is also illustrated in figure 7.5. There we evaluate numerically with the methods we exposed in this thesis the cumulative information on the amplitude of the ΛCDM power spectrum assuming lognormal statistics, for the one-dimensional configuration of points introduced in chapter 6, figure 6.1. The minimal distance we took is $r_{\rm min} = 10$ Mpc, corresponding to a variance of unity. On the other hand, the radius R plays no role. On the left panel is shown the contribution of the different wavenumbers, as a function of the maximal wavenumber extracted. We see clearly that, as expected, when entering the non linear regime the spectrum ceases slowly to contain further independent information. On the other hand, the same analysis with the real space correlation function, shown in the right panel is very different. There, almost the entirety of the information is contained in the variance, the two-point function at non-zero lag carrying very little information. In other words, for such parameters, there is as much information in the variance of the field than in its entire two-point function. We find this result an interesting starting point for further investigations, as this implies that it may be possible to understand the information content of the non linear regime with much simpler means than previously thought.

Certainly, a number of important issues still have to be answered. For instance, we have reached our conclusions using rather simple models such as the lognormal field. These models are good enough so that we can be sure that our claims on the N-point moment hierarchy of the noise-free fields do hold, at least for the first members of the hierarchy, which for practical reasons are the one that can be measured. However, when it comes to more stringent tests or measurements of a cosmological model, it is necessary to go beyond these simple prescriptions. Also, we pointed out on several occasions that underdense regions carry the largest part of the information in the nonlinear regime, a part unaccessible to the hierarchy. It is however not clear yet how this information can be extracted. As we discussed, in recent times several works have been looking at non linear transformations for the purpose of parameter inference, using expensive N-body simulations. Our simple analytical results bring strong support towards such methods, and provide as well an unambiguous explanation of their results, which are natural consequences of the statistics of fields with heavy tails. Nevertheless, it still remains to try and apply these methods to actual data. In light of our research, we judge this approach as promising, as it is a very straightforward manner to try and capture these large amount of information within

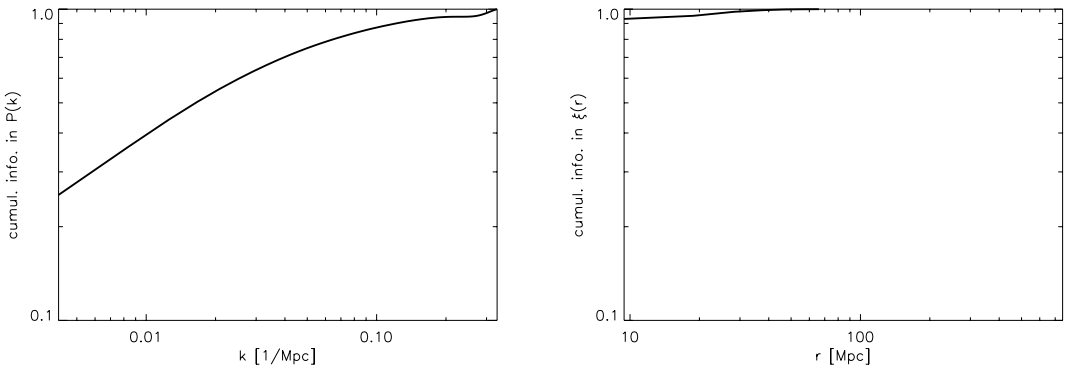

Fig. 7.5 — On the left panel it is shown the exact cumulative information within the spectrum of the one-dimensional lognormal field as function of the maximal wavenumber included in the analysis, for a variance of unity. The model parameter of interest is here the amplitude of the spectrum. For a Gaussian field, this line would be a straight line, but in the lognormal field saturation appears after the linear scales. On the right the same cumulative information but within the different arguments of the two-point function. The variance of the field carries the vast majority of the information.

the underdense regions. Clearly, the issue of the noise pervading the measurements will be essential for the success of these nonlinear transforms. We have tried to give very first elements of answers in this thesis, but a lot of work is still ahead. We hope that the methods introduced in this research will contribute to these efforts.

Bibliography

M. Abramowitz and I. Stegun. *Handbook of mathematical functions*. Dover Publications Inc., New York, 1970.

J. Aitchison and J.A. C Brown. *The lognormal distribution*. Cambridge University Press, 1957.

N.I. Akhiezer. *The classical moment problem and some related questions in analysis*. Oliver and Boyd, London, 1965.

A. Albrecht, G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox, J. C. Mather, S. Staggs, and N. B. Suntzeff. Report of the Dark Energy Task Force. *ArXiv Astrophysics e-prints*, September 2006.

A. Amara and A. Réfrégier. Optimal surveys for weak-lensing tomography. MNRAS, 381: 1018–1026, November 2007. doi: 10.1111/j.1365-2966.2007.12271.x.

L. Amendola. Non-Gaussian likelihood function. *Astrophysical Letters and Communications*, 33:63–69, February 1996.

G. E. Andrews, R. Askey, and R. Roy. *Special Functions*. Cambridge University Press, 1999.

D. J. Bacon, D. M. Goldberg, B. T. P. Rowe, and A. N. Taylor. Weak gravitational flexion. MNRAS, 365:414–428, January 2006. doi: 10.1111/j.1365-2966.2005.09624.x.

M. Bartelmann and P. Schneider. Weak gravitational lensing. Phys. Rep., 340:291–472, January 2001.

F. Bernardeau. The effects of smoothing on the statistical properties of large-scale cosmic fields. A&A, 291:697–712, November 1994.

F. Bernardeau and L. Kofman. Properties of the cosmological density distribution function. *The Astrophysical Journal*, 443:479–498, April 1995. doi: 10.1086/175542.

F. Bernardeau, S. Colombi, E. Gaztañaga, and R. Scoccimarro. Large-scale structure of the Universe and cosmological perturbation theory. Phys. Rep., 367:1–248, September 2002. doi: 10.1016/S0370-1573(02)00135-7.

G. M. Bernstein. Comprehensive Two-Point Analyses of Weak Gravitational Lensing Surveys. *The Astrophysical Journal*, 695:652–665, April 2009. doi: 10.1088/0004-637X/695/1/652.

Rajendra Bhatia. *Positive Definite Matrices*. Princeton Series in Applied Mathematics, 2007.

S. Blinnikov and R. Moessner. Expansions for nearly Gaussian distributions. A&AS, 130: 193–205, May 1998. doi: 10.1051/aas:1998221.

S. L. Bridle, M. P. Hobson, A. N. Lasenby, and R. Saunders. A maximum-entropy method for reconstructing the projected mass distribution of gravitational lenses. MNRAS, 299: 895–903, September 1998. doi: 10.1046/j.1365-8711.1998.01877.x.

R. Caldwell and M. Kamionkowski. Cosmology: Dark matter and dark energy. Nature, 458:587–589, April 2009. doi: 10.1038/458587a.

J. Carron. On the Incompleteness of the Moment and Correlation Function Hierarchy as Probes of the Lognormal Field. *The Astrophysical Journal*, 738:86–+, September 2011. doi: 10.1088/0004-637X/738/1/86.

J. Carron. On the assumption of Gaussian likelihoods for estimators of cosmological power spectra and their information content. *ArXiv e-prints*, April 2012a.

J. Carron. Information escaping the correlation hierarchy of the convergence field in the study of cosmological parameters. *Phys. Rev. Lett.*, 108:071301, Feb 2012b. doi: 10.1103/PhysRevLett.108.071301. URL `http://link.aps.org/doi/10.1103/PhysRevLett.108.071301`.

J. Carron and M. C. Neyrinck. On the Inadequacy of N-point Correlation Functions to Describe Nonlinear Cosmological Fields: Explicit Examples and Connection to Simulations. *The Astrophysical Journal*, 750:28, May 2012. doi: 10.1088/0004-637X/750/1/28.

J. Carron, A. Amara, and S. J. Lilly. Probe combination in large galaxy surveys: application of Fisher information and Shannon entropy to weak lensing. MNRAS, 417: 1938–1951, November 2011. doi: 10.1111/j.1365-2966.2011.19382.x.

A. Caticha. Lectures on Probability, Entropy, and Statistical Physics. *ArXiv e-prints*, July 2008.

M. Chevallier and D. Polarski. Accelerating Universes with Scaling Dark Matter. *International Journal of Modern Physics D*, 10:213–223, 2001. doi: 10.1142/S0218271801000822.

P. Coles and B. Jones. A lognormal model for the cosmological mass distribution. MNRAS, 248:1–13, January 1991.

S. Colombi. A 'skewed' lognormal approximation to the probablility distribution function of the large-scale density field. *The Astrophysical Journal*, 435:536–539, November 1994. doi: 10.1086/174834.

A. Cooray and R. Sheth. Halo models of large scale structure. Phys. Rep., 372:1–129, December 2002. doi: 10.1016/S0370-1573(02)00276-4.

S. Das and J. P. Ostriker. Testing a New Analytic Model for Gravitational Lensing Probabilities. *The Astrophysical Journal*, 645:1–16, July 2006. doi: 10.1086/504032.

S. Dodelson. *Modern cosmology*. Academic Press, 2003.

Charles F. Dunkl and Yuan Xu. *Orthogonal Polynomials of Several Variables*. Cambridge University Press, 2001.

R. Durrer. *The Cosmic Microwave Background*. Cambridge University Press, 2008.

T. Eifler, P. Schneider, and J. Hartlap. Dependence of cosmic shear covariances on cosmology. Impact on parameter estimation. A&A, 502:721–731, August 2009. doi: 10.1051/0004-6361/200811276.

D. J. Eisenstein and W. Hu. Baryonic Features in the Matter Transfer Function. *The Astrophysical Journal*, 496:605–+, March 1998. doi: 10.1086/305424.

X. Er, G. Li, and P. Schneider. Mass reconstruction by gravitational shear and flexion. *ArXiv e-prints*, August 2010.

R. A. Fisher. Theory of statistical estimation. *Proceedings of the Cambridge Philosophical Society*, 22:700–725, 1925.

Géza Freud. *Orthogonal Polynomials*. Pergamon Press Ltd., Headington Hill Hall, Oxford., 1971.

J. A. Frieman, M. S. Turner, and D. Huterer. Dark Energy and the Accelerating Universe. ARA&A, 46:385–432, September 2008. doi: 10.1146/annurev.astro.46.060407.145243.

J. N. Fry. Cosmological density fluctuations and large-scale structure From N-point correlation functions to the probability distribution. *The Astrophysical Journal*, 289:10–17, February 1985. doi: 10.1086/162859.

W. Gautschi. *Orthogonal Polynomials, Computation and Approximation*. Oxford University Press, 2004.

W. Gautschi. On euler's attempt to compute logarithms by interpolation: A commentary to his letter of february 16, 1734 to daniel bernoulli. *Journal of Computational and Applied Mathematics*, 219:408–415, 2008.

S. Hamimeche and A. Lewis. Likelihood analysis of CMB temperature and polarization power spectra. Phys. Rev. D, 77(10):103013, May 2008. doi: 10.1103/PhysRevD.77.103013.

A. Heavens. Fisher Matrices and All That: Experimental Design and Data Compression. In V. J. Martínez, E. Saar, E. Martínez-González, & M.-J. Pons-Bordería, editor, *Data Analysis in Cosmology*, volume 665 of *Lecture Notes in Physics, Berlin Springer Verlag*, pages 51–72, 2009. doi: 10.1007/978-3-540-44767-2_2.

K. Heitmann, M. White, C. Wagner, S. Habib, and D. Higdon. The Coyote Universe. I. Precision Determination of the Nonlinear Matter Power Spectrum. *The Astrophysical Journal*, 715:104–121, May 2010. doi: 10.1088/0004-637X/715/1/104.

C. C Heyde. On a property of the lognormal distribution. *The Journal of the Royal Statistical Society Series B (Methodological)*, 25(2):392–393, 1963.

W. Hu and B. Jain. Joint galaxy-lensing observables and the dark energy. Phys. Rev. D, 70(4):043009–+, August 2004. doi: 10.1103/PhysRevD.70.043009.

W. Hu and M. Tegmark. Weak Lensing: Prospects for Measuring Cosmological Parameters. ApJ, 514:L65–L68, April 1999. doi: 10.1086/311947.

R. G. Jarrett. Bounds and expansions for Fisher information when the moments are known. *Biometrika*, 71(1):101–113, 1984.

E. T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106(4):620–630, May 1957. doi: 10.1103/PhysRev.106.620. URL `http://prola.aps.org/abstract/PR/v106/i4/p620_1`.

E. T. Jaynes. *Papers On Probability, Statistics and Statistical Physics*. D. Reidel publishing Co., Dordrecht, Holland, 1983.

E. T. Jaynes and G. L. Bretthorst. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003. URL `http://omega.math.albany.edu:8008/JaynesBook.html`.

R. I. Jennrich and R. H Moore. Maximum likelihood estimation by means of nonlinear least squares. *Proceedings of the statistical computing section of the American statistical association*, pages 57–65, 1975.

G. Jungman, M. Kamionkowski, A. Kosowsky, and D. N. Spergel. Cosmological-parameter determination with microwave background maps. Phys. Rev. D, 54:1332–1344, July 1996a. doi: 10.1103/PhysRevD.54.1332.

G. Jungman, M. Kamionkowski, A. Kosowsky, and D. N. Spergel. Weighing the Universe with the Cosmic Microwave Background. *Physical Review Letters*, 76:1007–1010, February 1996b. doi: 10.1103/PhysRevLett.76.1007.

R. Juszkiewicz, D. H. Weinberg, P. Amsterdamski, M. Chodorowski, and F. Bouchet. Weakly nonlinear Gaussian fluctuations and the edgeworth expansion. *The Astrophysical Journal*, 442:39–56, March 1995. doi: 10.1086/175420.

J. Kac and P. Cheung. *Quantum calculus*. Springer, 2001.

N. Kaiser and G. Squires. Mapping the dark matter with weak gravitational lensing. *The Astrophysical Journal*, 404:441–450, February 1993. doi: 10.1086/172297.

I. Kayo, A. Taruya, and Y. Suto. Probability Distribution Function of Cosmological Density Fluctuations from a Gaussian Initial Condition: Comparison of One-Point and Two-Point Lognormal Model Predictions with N-Body Simulations. *The Astrophysical Journal*, 561:22–34, November 2001. doi: 10.1086/323227.

E. Koelink and W. Van Assche. Leonhard Euler and a q-analogue of the logarithm. *Proceedings of The American Mathematical Society*, 137(05):1663–1676, 2009.

E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, D. N. Spergel, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. ApJS, 192:18, February 2011. doi: 10.1088/0067-0049/192/2/18.

B. O. Kopman. On distributions admitting a sufficient statistic. *Trans. Am. Math. Soc.*, 39(3):399–509, 1936.

E. Lawrence, K. Heitmann, M. White, D. Higdon, C. Wagner, S. Habib, and B. Williams. The Coyote Universe. III. Simulation Suite and Precision Emulator for the Nonlinear Matter Power Spectrum. *The Astrophysical Journal*, 713:1322–1331, April 2010. doi: 10.1088/0004-637X/713/2/1322.

J. Lee and U.-L. Pen. Information Content in the Galaxy Angular Power Spectrum from the Sloan Digital Sky Survey and Its Implication on Weak-Lensing Analysis. ApJ, 686: L1–L4, October 2008. doi: 10.1086/592820.

A. R. Liddle and D. H. Lyth, editors. *Cosmological inflation and large-scale structure*. New York : Cambridge University Press, 2000., 2000.

E. V. Linder. Exploring the Expansion History of the Universe. *Physical Review Letters*, 90(9):091301–+, March 2003.

K. Maisinger, M. P. Hobson, and A. N. Lasenby. Maximum-entropy image reconstruction using wavelets. MNRAS, 347:339–354, January 2004. doi: 10.1111/j.1365-2966.2004.07216.x.

P. J. Marshall, M. P. Hobson, S. F. Gull, and S. L. Bridle. Maximum-entropy weak lens reconstruction: improved methods and application to data. MNRAS, 335:1037–1048, October 2002. doi: 10.1046/j.1365-8711.2002.05685.x.

R. Massey, B. Rowe, A. Refregier, D. J. Bacon, and J. Bergé. Weak gravitational shear and flexion with polar shapelets. MNRAS, 380:229–245, September 2007. doi: 10.1111/j.1365-2966.2007.12072.x.

T. Matsubara. Nonlinear perturbation theory integrated with nonlocal bias, redshift-space distortions, and primordial non-Gaussianity. Phys. Rev. D, 83(8):083518–+, April 2011. doi: 10.1103/PhysRevD.83.083518.

H. Mo, F. Van Den Bosch, and S.D.M. White. *Galaxy Formation and Evolution*. Galaxy Formation and Evolution. Cambridge University Press, 2010. ISBN 9780521857932. URL `http://books.google.ch/books?id=Zj7fDU3Z4wsC`.

D. Munshi, P. Valageas, L. Van Waerbeke, and A. Heavens. Cosmology with Weak Lensing Surveys. *ArXiv Astrophysics e-prints*, December 2006.

M. C. Neyrinck. Rejuvenating the Matter Power Spectrum. III. The Cosmology Sensitivity of Gaussianized Power Spectra. *The Astrophysical Journal*, 742:91, December 2011. doi: 10.1088/0004-637X/742/2/91.

M. C. Neyrinck, I. Szapudi, and C. D. Rimes. Information content in the halo-model dark-matter power spectrum. MNRAS, 370:L66–L70, July 2006. doi: 10.1111/j.1745-3933.2006.00190.x.

M. C. Neyrinck, I. Szapudi, and A. S. Szalay. Rejuvenating the Matter Power Spectrum: Restoring Information with a Logarithmic Density Mapping. ApJ, 698:L90–L93, June 2009. doi: 10.1088/0004-637X/698/2/L90.

M. C. Neyrinck, I. Szapudi, and A. S. Szalay. Rejuvenating Power Spectra. II. The Gaussianized Galaxy Density Field. *The Astrophysical Journal*, 731:116–+, April 2011. doi: 10.1088/0004-637X/731/2/116.

Y. Okura, K. Umetsu, and T. Futamase. A New Measure for Weak-Lensing Flexion. *The Astrophysical Journal*, 660:995–1002, May 2007. doi: 10.1086/513135.

D. Parkinson, C. Blake, M. Kunz, B. A. Bassett, R. C. Nichol, and K. Glazebrook. Optimizing baryon acoustic oscillation surveys - I. Testing the concordance ΛCDM cosmology. MNRAS, 377:185–197, May 2007. doi: 10.1111/j.1365-2966.2007.11586.x.

P. J. E. Peebles. *The large-scale structure of the universe*. Princeton University Press, 1980.

W. J. Percival, et al, and . Baryon acoustic oscillations in the Sloan Digital Sky Survey Data Release 7 galaxy sample. MNRAS, 401:2148–2168, February 2010. doi: 10.1111/j.1365-2966.2009.15812.x.

A. Pillepich, C. Porciani, and O. Hahn. Halo mass function and scale-dependent bias from N-body simulations with non-Gaussian initial conditions. MNRAS, 402:191–206, February 2010. doi: 10.1111/j.1365-2966.2009.15914.x.

S. Pires and A. Amara. Weak Lensing Mass Reconstruction: Flexion Versus Shear. *The Astrophysical Journal*, 723:1507–1511, November 2010. doi: 10.1088/0004-637X/723/2/1507.

E. J. G. Pitman and J. Wishart. Sufficient statistics and intrinsic accuracy. In *Proceedings of the Cambridge Philosophical Society*, volume 32 of *Proceedings of the Cambridge Philosophical Society*, pages 567–+, 1936. doi: 10.1017/S0305004100019307.

A. C. Pope, T. Matsubara, A. S. Szalay, M. R. Blanton, D. J. Eisenstein, J. Gray, B. Jain, N. A. Bahcall, J. Brinkmann, T. Budavari, A. J. Connolly, J. A. Frieman, J. E. Gunn, D. Johnston, S. M. Kent, R. H. Lupton, A. Meiksin, R. C. Nichol, D. P. Schneider, R. Scranton, M. A. Strauss, I. Szapudi, M. Tegmark, M. S. Vogeley, D. H. Weinberg, I. Zehavi, and SDSS Collaboration. Cosmological Parameters from Eigenmode Analysis of Sloan Digital Sky Survey Galaxy Redshifts. *The Astrophysical Journal*, 607:655–660, June 2004. doi: 10.1086/383533.

C.Radhkrishna Rao. *Lineare statistische Methoden und ihre Anwendungen*. Akademie Verlag Berlin, 1973.

A. Refregier. Weak Gravitational Lensing by Large-Scale Structure. ARA&A, 41:645–668, 2003.

A. Refregier, A. Amara, T. Kitching, and A. Rassat. iCosmo: an Interactive Cosmology Package. *ArXiv e-prints*, October 2008.

M. Riesz. Sur le problème des moments et le théorème de parseval correspondant. *Acta Syeged Sect. Math.*, 1:209–225, 1923.

C. D. Rimes and A. J. S. Hamilton. Information content of the non-linear matter power spectrum. MNRAS, 360:L82–L86, June 2005. doi: 10.1111/j.1745-3933.2005.00051.x.

P. Schneider and X. Er. Weak lensing goes bananas: what flexion really measures. A&A, 485:363–376, July 2008. doi: 10.1051/0004-6361:20078631.

P. Schneider, J. Ehlers, and E. E. Falco. *Gravitational Lenses*. Gravitational Lenses, XIV, 560 pp. 112 figs.. Springer-Verlag Berlin Heidelberg New York. Also Astronomy and Astrophysics Library, 1992.

P. Schneider, C. S. Kochanek, and J. Wambsganss. *Gravitational Lensing: Strong, Weak and Micro*. 2006. doi: 10.1007/978-3-540-30310-7.

S. Seitz and P. Schneider. A new finite-field mass reconstruction algorithm. A&A, 374: 740–745, August 2001. doi: 10.1051/0004-6361:20010493.

H.-J. Seo and D. J. Eisenstein. Probing Dark Energy with Baryonic Acoustic Oscillations from Future Large Galaxy Redshift Surveys. *The Astrophysical Journal*, 598:720–740, December 2003. doi: 10.1086/379122.

H.-J. Seo and D. J. Eisenstein. Improved Forecasts for the Baryon Acoustic Oscillations and Cosmological Distance Scale. *The Astrophysical Journal*, 665:14–24, August 2007. doi: 10.1086/519549.

H.-J. Seo, M. Sato, S. Dodelson, B. Jain, and M. Takada. Re-capturing Cosmic Information. ApJ, 729:L11+, March 2011. doi: 10.1088/2041-8205/729/1/L11.

C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423,623–656, 1948.

R. Shohat and J. D Tamarkin. *The Problem of Moments*. American Mathematical Society, third printing of revised edition, 1963.

B. Simon. The Christoffel-Darboux Kernel. *ArXiv e-prints*, June 2008.

Barry Simon. The classical moment problem as a self-adjoint finite difference operator. *Adv. Math*, 137:82–203, 1997.

J. Skilling and R. K. Bryan. Maximum Entropy Image Reconstruction - General Algorithm. MNRAS, 211:111–+, November 1984.

R. E. Smith, J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce, P. A. Thomas, G. Efstathiou, and H. M. P. Couchman. Stable clustering, the halo model and non-linear cosmological power spectra. MNRAS, 341:1311–1332, June 2003.

J.-L. Starck and E. Pantin. Multiscale maximum entropy images restoration. *Vistas in Astronomy*, 40:563–569, 1996. doi: 10.1016/S0083-6656(96)00042-6.

T. J. Stieltjes. Recherches sur les fractions continues. *Annales de la faculté des sciences de Toulouse Sér. 1*, 8(4):J1–J122, 1894.

J. Stoyanov. *Counterexamples in probability*. Wiley, 1987. ISBN 9780471916499.

J. Stoyanov. Stieltjes classes for moment-indeterminate probability distributions. *Journal of Applied Probability*, 41A:281–294, 2004.

I. Szapudi. Introduction to Higher Order Spatial Statistics in Cosmology. *ArXiv Astrophysics e-prints*, May 2005.

G. Szegö. *Orthogonal Polynomials*. American Mathematical Society, fourth edition, 2003.

R. Takahashi, M. Oguri, M. Sato, and T. Hamana. Probability Distribution Functions of Cosmological Lensing: Convergence, Shear, and Magnification. *The Astrophysical Journal*, 742:15, November 2011. doi: 10.1088/0004-637X/742/1/15.

A. Taruya, M. Takada, T. Hamana, I. Kayo, and T. Futamase. Lognormal Property of Weak-Lensing Fields. *The Astrophysical Journal*, 571:638–653, June 2002. doi: 10.1086/340048.

A. N. Taylor and P. I. R. Watts. Evolution of the cosmological density distribution function. MNRAS, 314:92–98, May 2000. doi: 10.1046/j.1365-8711.2000.03339.x.

A. N. Taylor and P. I. R. Watts. Parameter information from non-linear cosmological fields. MNRAS, 328:1027–1038, December 2001. doi: 10.1046/j.1365-8711.2001.04874.x.

M. Tegmark. Measuring Cosmological Parameters with Galaxy Surveys. *Physical Review Letters*, 79:3806–3809, November 1997. doi: 10.1103/PhysRevLett.79.3806.

M. Tegmark, A. N. Taylor, and A. F. Heavens. Karhunen-Loeve Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets? *The Astrophysical Journal*, 480:22–+, May 1997. doi: 10.1086/303939.

M. Tegmark, D. J. Eisenstein, M. A. Strauss, D. H. Weinberg, M. R. Blanton, J. A. Frieman, M. Fukugita, J. E. Gunn, A. J. S. Hamilton, G. R. Knapp, R. C. Nichol, J. P. Ostriker, N. Padmanabhan, W. J. Percival, D. J. Schlegel, D. P. Schneider, R. Scoccimarro, U. Seljak, H.-J. Seo, M. Swanson, A. S. Szalay, M. S. Vogeley, J. Yoo, I. Zehavi, K. Abazajian, S. F. Anderson, J. Annis, N. A. Bahcall, B. Bassett, A. Berlind, J. Brinkmann, T. Budavari, F. Castander, A. Connolly, I. Csabai, M. Doi, D. P. Finkbeiner, B. Gillespie, K. Glazebrook, G. S. Hennessy, D. W. Hogg, Ž. Ivezić, B. Jain, D. Johnston, S. Kent, D. Q. Lamb, B. C. Lee, H. Lin, J. Loveday, R. H. Lupton, J. A. Munn, K. Pan, C. Park, J. Peoples, J. R. Pier, A. Pope, M. Richmond, C. Rockosi, R. Scranton, R. K. Sheth, A. Stebbins, C. Stoughton, I. Szapudi, D. L. Tucker, D. E. vanden Berk, B. Yanny, and D. G. York. Cosmological constraints from the SDSS luminous red galaxies. Phys. Rev. D, 74(12):123507, December 2006. doi: 10.1103/PhysRevD.74.123507.

The Dark Energy Survey Collaboration. The Dark Energy Survey. *ArXiv Astrophysics e-prints*, October 2005.

Evgenij E. Tyrtyshnikov. How bad are hankel matrices? *Numerische Mathematik*, 67: 261–269, 1994. doi: 10.1007/s002110050027.

A. Vallinotto, S. Dodelson, and P. Zhang. Magnification as a Tool in Weak Lensing. *ArXiv e-prints*, September 2010.

A. van den Bos. *Parameter estimation for scientists and engineers*. John Wiley and Sons, NY, 2007.

M. Velander, K. Kuijken, and T. Schrabback. Probing Galaxy Dark Matter Haloes in COSMOS with Weak Lensing Flexion. *ArXiv e-prints*, November 2010.

M. S. Vogeley and A. S. Szalay. Eigenmode Analysis of Galaxy Redshift Surveys. I. Theory and Methods. *The Astrophysical Journal*, 465:34, July 1996. doi: 10.1086/177399.

X. Wang, M. Neyrinck, I. Szapudi, A. Szalay, X. Chen, J. Lesgourgues, A. Riotto, and M. Sloth. Perturbation Theory of the Cosmological Log-density Field. *The Astrophysical Journal*, 735:32–+, July 2011. doi: 10.1088/0004-637X/735/1/32.

S. Wigert. Sur les polynomes orthogonaux et l'approximation des functions continues. *Arkiv för matematik, astronomi och fysik*, 17(18), 1923.

V. Wild, J. A. Peacock, O. Lahav, E. Conway, S. Maddox, I. K. Baldry, C. M. Baugh, J. Bland-Hawthorn, T. Bridges, R. Cannon, S. Cole, M. Colless, C. Collins, W. Couch, G. Dalton, R. De Propris, S. P. Driver, G. Efstathiou, R. S. Ellis, C. S. Frenk, K. Glazebrook, C. Jackson, I. Lewis, S. Lumsden, D. Madgwick, P. Norberg, B. A. Peterson, W. Sutherland, and K. Taylor. The 2dF Galaxy Redshift Survey: stochastic relative biasing between galaxy populations. MNRAS, 356:247–269, January 2005. doi: 10.1111/j.1365-2966.2004.08447.x.

Herbert S. Wilf. *Generatingfunctionology*. Academic Press, Inc., 2 edition, 1994.

L. Wolz, M. Kilbinger, J. Weller, and T. Giannantonio. On the Validity of Cosmological Fisher Matrix Forecasts. *ArXiv e-prints*, May 2012.

Konstantinos Zografos and Kosmas Ferentinos. An information theoretic argument for the validity of the exponential model. *Metrika*, 41(1):109–119, December 1994. URL `http://ideas.repec.org/a/spr/metrik/v41y1994i1p109-119.html`.

C. Zunckel and R. Trotta. Reconstructing the history of dark energy using maximum entropy. MNRAS, 380:865–876, September 2007. doi: 10.1111/j.1365-2966.2007.12000.x.